

Maximal Sharing in the Lambda Calculus with letrec¹

Clemens Grabmayer

Dept. of Computer Science, VU University Amsterdam
de Boelelaan 1081a, 1081 HV Amsterdam
c.a.grabmayer@vu.nl

Jan Rochel

Dept. of Computing Sciences, Utrecht University
Princetonplein 5, 3584 CC Utrecht, The Netherlands
jan@rochel.info

Abstract

Increasing sharing in programs is desirable to compactify the code, and to avoid duplication of reduction work at run-time, thereby speeding up execution. We show how a maximal degree of sharing can be obtained for programs expressed as terms in the lambda calculus with letrec. We introduce a notion of ‘maximal compactness’ for λ_{letrec} -terms among all terms with the same infinite unfolding. Instead of defined purely syntactically, this notion is based on a graph semantics. λ_{letrec} -terms are interpreted as first-order term graphs so that unfolding equivalence between terms is preserved and reflected through bisimilarity of the term graph interpretations. Compactness of the term graphs can then be compared via functional bisimulation.

We describe practical and efficient methods for the following two problems: transforming a λ_{letrec} -term into a maximally compact form; and deciding whether two λ_{letrec} -terms are unfolding-equivalent. The transformation of a λ_{letrec} -term L into maximally compact form L_0 proceeds in three steps: (i) translate L into its term graph $G = \llbracket L \rrbracket$; (ii) compute the maximally shared form of G as its bisimulation collapse G_0 ; (iii) read back a λ_{letrec} -term L_0 from the term graph G_0 with the property $\llbracket L_0 \rrbracket = G_0$. This guarantees that L_0 and L have the same unfolding, and that L_0 exhibits maximal sharing.

The procedure for deciding whether two given λ_{letrec} -terms L_1 and L_2 are unfolding-equivalent computes their term graph interpretations $\llbracket L_1 \rrbracket$ and $\llbracket L_2 \rrbracket$, and checks whether these term graphs are bisimilar.

For illustration, we also provide a readily usable implementation.

Categories and Subject Descriptors D.3.3 [Language constructs and features]: Recursion; F.3.3 [Studies of Programming Constructs]: Functional constructs

General Terms functional programming, compiler optimisation

Keywords Lambda Calculus with letrec, unfolding semantics, subterm sharing, maximal sharing, higher-order term graphs

1. Introduction

Explicit sharing in pure functional programming languages is typically expressed by means of the letrec construct, which facilitates cyclic definitions. The λ -calculus with letrec, λ_{letrec} forms a syntactic core of these languages, and it can be viewed as their abstraction. As such λ_{letrec} is well-suited as a test bed for developing program transformations in functional programming languages. This certainly holds for the transformation presented here that has a strong conceptual motivation, is justified by a form of semantic reasoning, and is best described first for an expressive, yet minimal language.

¹This work was supported by NWO in the framework of the project *Realising Optimal Sharing (ROS)*, project number 612.000.935.

1.1 Expressing sharing and infinite λ -terms

For the programmer the letrec-construct offers the possibility to write a program more compactly by utilising subterm sharing. letrec-expressions bind subterms to variables; these variables then denote occurrences of the respective subterms and can be used anywhere inside of the letrec-expression (also recursively). In this way, instead of repeating a subterm multiple times, a single definition can be given which is then referenced from multiple positions.

We will denote the letrec-construct here by let as in Haskell.

Example 1.1. Consider the λ -term $(\lambda x. x) (\lambda x. x)$ with two occurrences of the subterm $\lambda x. x$. These occurrences can be shared with as result the λ_{letrec} -term $(\text{let } id = \lambda x. x \text{ in } id \ id)$.

As let-bindings permit definitions with cyclic dependencies, terms in λ_{letrec} are able to finitely denote infinite λ -terms (for short: λ^∞ -terms). The λ^∞ -term M represented by a λ_{letrec} -term L can be obtained by a typically infinite process in which the let-bindings in L are unfolded continually with M as result in the limit. Then we say that M is the *infinite unfolding* of L , or that M is the denotation of L in the *unfolding semantics*, indicated symbolically by $M = \llbracket L \rrbracket_{\lambda^\infty}$.

Example 1.2. For the λ_{letrec} -terms L and P and the λ^∞ -term M :

$$\begin{aligned} L &:= \lambda f. \text{let } r = f \ r \text{ in } r & M &:= \lambda f. f \ (f \ (\dots)) \\ P &:= \lambda f. \text{let } r = f \ (f \ r) \text{ in } r \end{aligned}$$

it holds that both L and P (which represent fixed-point combinators) have M as their infinite unfolding: $\llbracket L \rrbracket_{\lambda^\infty} = \llbracket P \rrbracket_{\lambda^\infty} = M$.

L and P in this example are ‘unfolding equivalent’. Note that L represents M in a more compact way than P . It is intuitively clear that there is no λ_{letrec} -term that represents M more compactly than L . So L can be called a ‘maximally shared form’ of P (and of M).

We address, and efficiently solve, the problems of computing the maximally shared form of a λ_{letrec} -term, and of determining whether two λ_{letrec} -terms are unfolding-equivalent. Note that these notions are based on the static unfolding semantics. We *do not consider* any dynamic semantics based on evaluation by β -reduction or otherwise.

1.2 Recognising potential for sharing

A general risk for compilers of functional programs is “[to construct] multiple instances of the same expression, rather than sharing a single copy of them. This wastes space because each instance occupies separate storage, and it wastes time because the instances will be reduced separately. This waste can be arbitrarily large, [...]” ([29, p.243]). Therefore practical compilers increase sharing, and do so typically for supercombinator translations of programs (such as fully-lazy lambda-lifting). Thereby two goals are addressed: to increase sharing based on a syntactical analysis of the ‘static’ form of the program; and to prevent splits into too many supercombinators

when an anticipation of the program's 'dynamic' behaviour is able to conclude that no sharing at run-time will be gained.

A well-known method for the 'static' part is common sub-expression elimination (CSE) [6]. For the 'dynamic' part, a predictive syntactic program analysis has been proposed for fine-tuning sharing of partial applications in supercombinator translations [10].

We focus primarily on the 'static' aspect of introducing sharing. We provide a conceptual solution that substantially extends CSE. But instead of maximising sharing for a supercombinator translation of a program, we carry out the optimisation on the program itself (the λ_{letrec} -term). And instead of applying a purely syntactical program analysis, we use a term graph semantics for λ_{letrec} -terms.

1.3 Approach based on a term graph semantics

We develop a combination of techniques for realising maximal sharing in λ_{letrec} -terms. For this we proceed in four steps: λ_{letrec} -terms are interpreted as higher-order term graphs; the higher-order term graphs are implemented as first-order term graphs; maximally compact versions of such term graphs can be computed by standard algorithms; λ_{letrec} -terms that represent compacted term graphs (or in fact arbitrary ones) can be retrieved by a 'readback' operation.

In more detail, the four essential ingredients are the following:

- (1) A *semantics* $\llbracket \cdot \rrbracket_{\mathcal{H}}$ for interpreting λ_{letrec} -terms as *higher-order term graphs*, which are first-order term graphs enriched with a feature for describing binding and scopes. We call this specific kind of higher-order term graphs ' λ -ho-term-graphs'.

The variable binding structure is recorded in this term graph concept because it must be respected by any addition of sharing. The term graph interpretation adequately represents sharing as expressed by a λ_{letrec} -term. It is not injective: a λ -ho-term-graph typically is the interpretation of various λ_{letrec} -terms. Different degrees of sharing as expressed by λ_{letrec} -terms can be compared via the λ -ho-term-graph interpretations by a sharing preorder, which is defined as the existence of a homomorphism (functional bisimulation).

While comparing higher-order term graphs via this preorder is computable in principle, standard algorithms do not apply. Therefore efficient solvability of the compactification problem and the comparison problem is, from the outset, not guaranteed. For this reason we devise a first-order implementation of λ -ho-term-graphs:

- (2) An *interpretation* \mathcal{HT} of λ -ho-term-graphs into a specific kind of *first-order term graphs*, which we call ' λ -term-graphs'. It preserves and reflects the sharing preorder.

\mathcal{HT} reduces bisimilarity between λ -ho-term-graphs (higher-order) to bisimilarity between λ -term-graphs (first-order), and facilitates:

- (3) The use of standard methods for *checking* bisimilarity and for computing the bisimulation *collapse* of λ -term-graphs. Via \mathcal{HT} also the analogous problems for λ -ho-term-graphs can be solved.

Term graphs can be represented as deterministic process graphs (labelled transition systems), and even as deterministic finite-state automata (DFAs). That is why it is possible to apply efficient algorithms for state minimisation and language equivalence of DFAs.

Finally, an operation to return from term graphs to λ_{letrec} -terms:

- (4) A *readback* function rb from λ -term-graphs to λ_{letrec} -terms that, for every λ -term-graph G , computes a λ_{letrec} -term L from the set of λ_{letrec} -terms that have G as their interpretation via $\llbracket \cdot \rrbracket_{\mathcal{H}}$ and \mathcal{HT} (i.e. a λ_{letrec} -term for which it holds that $\mathcal{HT}(\llbracket L \rrbracket_{\mathcal{H}}) = G$).

1.4 Methods and their correctness

On the basis of the concepts above we develop efficient methods for introducing maximal sharing, and for checking unfolding equivalence, of λ_{letrec} -terms, as sketched below.

In describing these methods, we use the following notation:




Figure 1. Component-step build-up of the methods for computing a maximally shared form L_0 of a λ_{letrec} -term L (left), and for deciding unfolding equivalence of λ_{letrec} -terms L_1 and L_2 via bisimilarity \Leftrightarrow (right).




Figure 2. Computing a maximally compact version of the term P from Ex. 1.2 (right) by using composition of term graph semantics $\llbracket \cdot \rrbracket_{\mathcal{T}}$, collapse $| \downarrow$, and readback rb , yielding the term L (left).

- \mathcal{H} : class of λ -ho-term-graphs, the image of the semantics $\llbracket \cdot \rrbracket_{\mathcal{H}}$;
- \mathcal{T} : class of λ -term-graphs, the image of the interpretation \mathcal{HT} ;
- $\llbracket \cdot \rrbracket_{\mathcal{T}} := \mathcal{HT} \circ \llbracket \cdot \rrbracket_{\mathcal{H}}$: first-order term graph semantics for λ_{letrec} -terms;
- $| \downarrow$: bisimulation collapse on \mathcal{H} and \mathcal{T} ;
- rb : readback mapping from λ -term-graphs to λ_{letrec} -terms.

We obtain the following methods (for illustrations, see Fig. 1):

- ▷ *Maximal sharing*: for a given λ_{letrec} -term, a maximally shared form can be obtained by collapsing its first-order term graph interpretation, and then reading back the collapse: $\text{rb} \circ | \downarrow \circ \llbracket \cdot \rrbracket_{\mathcal{T}}$
- ▷ *Unfolding equivalence*: for given λ_{letrec} -terms L and P , it can be decided whether $\llbracket L \rrbracket_{\lambda^\infty} = \llbracket P \rrbracket_{\lambda^\infty}$ by checking whether their term graph interpretations $\llbracket L \rrbracket_{\mathcal{T}}$ and $\llbracket P \rrbracket_{\mathcal{T}}$ are bisimilar.

See Fig. 2 for an illustration of the application of the maximal sharing method to the λ_{letrec} -terms L and P from Ex. 1.2.

The correctness of these methods hinges on the fact that the term graph translation and the readback satisfy the following properties:

- (P1) λ_{letrec} -terms L and P have the same infinite unfolding if and only if the term graphs $\llbracket L \rrbracket_{\mathcal{T}}$ and $\llbracket P \rrbracket_{\mathcal{T}}$ are bisimilar.
- (P2) The class \mathcal{T} of λ -term-graphs is closed under homomorphism.
- (P3) The readback rb is a right inverse of $\llbracket \cdot \rrbracket_{\mathcal{T}}$ up to isomorphism \simeq , that is, for all term graphs $G \in \mathcal{T}$ it holds: $(\llbracket \cdot \rrbracket_{\mathcal{T}} \circ \text{rb})(G) \simeq G$.

Note: (P2) and (P3) will be established only for a subclass \mathcal{T}_{eag} of \mathcal{T} . Furthermore, practicality of these methods depends on the property:

- (P4) Translation $\llbracket \cdot \rrbracket_{\mathcal{T}}$ and readback rb are efficiently computable.

1.5 Overview of the development

In the Preliminaries (Section 2) we fix basic notions and notations for first-order term graphs. λ_{letrec} -terms and their unfolding semantics are defined in Section 3. In Section 4 we develop the concept of ‘ λ -ho-term-graph’, which gives rise to the class \mathcal{H} , and the higher-order term graph semantics $\llbracket \cdot \rrbracket_{\mathcal{H}}$ for λ_{letrec} -terms.

In Section 5 we develop the concept of first-order ‘ λ -term-graph’ in the class \mathcal{T} , and define the interpretation \mathcal{HT} of λ -ho-term-graphs into λ -term-graphs as a mapping from \mathcal{H} to \mathcal{T} . This induces the first-order term graph semantics $\llbracket \cdot \rrbracket_{\mathcal{T}} := \mathcal{HT} \circ \llbracket \cdot \rrbracket_{\mathcal{H}}$, for which we also provide a direct inductive definition.

In Section 6 we define the readback rb with the desired property as a function from λ -term-graphs to λ_{letrec} -terms. Subsequently in Section 7 we report on the complexity of the described methods, individually, and in total for the methods described in Subsection 1.4.

In Section 8 we link to our implementation of the presented methods. Finally in Section 9 we explain easy modifications, describe possible extensions, and sketch potential practical applications.

1.6 Applications and scalability

While our contribution is at first a conceptual one, it holds the promise for a number of practical applications:

- Increasing the efficiency of the execution of programs by transforming them into their maximally shared form at compile-time.
- Increasing the efficiency of the execution of programs by repeatedly compactifying the program at run time.
- Improving systems for recognising program equivalence.
- Providing feedback to the programmer, along the lines: ‘This code has identical fragments and can be written more compactly.’

These and a number of other potential applications are discussed in more detail in Section 9.

The presented methods scale well to larger inputs, due to the quadratic bound on their runtime complexity (see Section 7).

1.7 Relationship with other concepts of sharing

The maximal sharing method is targeted at increasing ‘static’ sharing: in the sense that a program is transformed at compile time into a version with a higher degree of sharing. It is not (at least not a priori) a method for ‘dynamic’ sharing, i.e. for an evaluator that maintains a certain degree of sharing at run time, such as graph rewrite mechanisms for fully-lazy [33] or optimal evaluation [1] of the λ -calculus. However, we envisage run-time collapsing of the program’s graph interpretation integrated with the evaluator (see Section 9).

The term ‘maximal sharing’ stems from work on the ATERM library [5]. It describes a technique for minimising memory usage when representing a set of terms in a first-order term rewrite system (TRS). The terms are kept in an aggregate directed acyclic graph by which their syntax trees are shared as much as possible. Thereby terms are created only if they are entirely new; otherwise they are referenced by pointers to roots of sub-dags. Our use of the expression ‘maximal sharing’ is inspired by that work, but our results generalise that approach in the following ways:

- Instead of first-order terms we consider terms in a higher-order language with the letrec -construct for expressing sharing.
- Since letrec typically defines cyclic sharing dependencies, we interpret terms as cyclic graphs instead of just dags.
- We are interested in increasing sharing by bisimulation collapse instead of by identifying isomorphic sub-dags.

ATERM only checks for equality of subexpressions. Therefore it only introduces horizontal sharing and implements a form of *common subexpression elimination (CSE)* [29, p. 241]. Our approach

is stronger than CSE: while Ex. 1.1 can be handled by CSE, this is not the case for Ex. 1.2. In contrast to CSE, our approach increases also vertical and twisted sharing². (see also [4]).

1.8 Contribution of this paper in context

Blom introduces *higher-order term graphs* [4], which are extensions of first-order term graphs by adding a scope function that assigns a set of vertices, its *scope*, to every abstraction vertex.

As a stepping stone for the methods we develop here, we use concepts and results that we described in an earlier paper [12]. There, for interpreting λ_{letrec} -terms, a modification of Blom’s higher-order term graphs (the λ -ho-term-graphs of the class \mathcal{H}) in which scopes are represented by means of ‘abstraction prefix functions’. We also investigated first-order λ -term-graphs with scope-delimiter vertices (corresponding to the class \mathcal{T} here). In particular we examined which specific class of first-order λ -term-graphs can faithfully represent the higher-order λ -ho-term-graphs in such a way that compactification of the latter can be realised through bisimulation collapse of the former (this led to the λ -term-graphs of the class \mathcal{T}).

Whereas in the paper [12] we exclusively focused on the graph formalisms, and investigated them in their own right, here we connect the results obtained there to the language λ_{letrec} for expressing sharing and cyclicity. Since the methods presented here are based on the graph formalisms, and rely on their properties for correctness, we recapitulate the concepts and the relevant results in Sec. 4 and 5.

The translation $\llbracket \cdot \rrbracket_{\mathcal{T}}$ of λ_{letrec} -terms into first-order term graphs was inspired by related representations that use scope delimiters to indicate end of scopes. Such representations are generalisations of a de Bruijn index notation for λ -terms [8] in which the de Bruijn indexes are numerals of the form $S(\dots(S(0))\dots)$. In the generalised form, due to Patterson and Bird [3], the symbol S can occur anywhere between a variable occurrence and its binding abstraction. The idea to view S as a scope delimiter was employed by Hendriks and van Oostrom, who defined an end-of-scope symbol \wedge [18]. This approach is also used in the translation of pure λ -terms (without letrec) into Lambdascope -graphs (interaction nets) on which van Oostrom defines an optimal evaluator for the λ -calculus [27].

We have also used these first-order representations of λ -terms with scope delimiters for studying the limits of an optimising program transformation that, for a given λ_{letrec} -term, contracts directly visible redexes, and, whenever possible, also contracts such redexes that are concealed by recursion [31]. The result of the optimisation should again be a λ_{letrec} -term. Since this program transformation can best be defined, for a given λ_{letrec} -term L , on the infinite λ -term M that is the unfolding semantics (see Section 3) of L , it is crucial to know when the result of contracting a development of redexes in M that corresponds to a visible or a concealed redex in L can again be written as a λ_{letrec} -terms.

This suggested the question: how can those infinite λ -terms be characterised that are expressible by λ_{letrec} -terms in the sense that they arise as the unfolding semantics of a λ_{letrec} -term? We answered this question for λ_{letrec} -terms in the report [11], and, obtaining the same answer, for λ_{μ} -terms in the article [13] with accompanying report [14]. We defined a rewrite system that decomposes λ -terms by steps that ‘observe’ λ -abstractions, applications, variable occurrences, and end of scopes. We showed that infinite λ -terms that are the unfolding semantics of λ_{letrec} -terms or of λ_{μ} -terms are precisely those that have only finitely many ‘generated subterms’, that is, reducts in the decomposition rewrite system.

2. Preliminaries

By \mathbb{N} we denote the natural numbers including zero. For words w over an alphabet A , the length of w is denoted by $|w|$.

²For definitions of horizontal, vertical, and twisted sharing we refer to [4].

Let Σ be a TRS-signature [32] with arity function $ar : \Sigma \rightarrow \mathbb{N}$. A *term graph* over Σ (or a Σ -*term-graph*) is a tuple $\langle V, lab, args, r \rangle$ where: V is a set of *vertices*, $lab : V \rightarrow \Sigma$ the (*vertex*) *label function*, $args : V \rightarrow V^*$ the *argument function* that maps every vertex v to the word $args(v)$ consisting of the $ar(lab(v))$ successor vertices of v (hence $|args(v)| = ar(lab(v))$), and r , the *root*, is a vertex in V . Term graphs may have infinitely many vertices.

Let G be a term graph over signature Σ . As useful notation for picking out an arbitrary vertex, or the i -th vertex, from among the ordered successors of a vertex v in G , we define for each $i \in \mathbb{N}$ the indexed edge relation $\rightarrow_i \subseteq V \times V$, and additionally the (not indexed) edge relation $\rightarrow \subseteq V \times V$, by stipulating for all $w, w' \in V$:

$$\begin{aligned} w \rightarrow_i w' &: \Leftrightarrow \exists w_0, \dots, w_n \in V. args(w) = w_0 \dots w_n \wedge w' = w_i \\ w \rightarrow w' &: \Leftrightarrow \exists i \in \mathbb{N}. w \rightarrow_i w' \end{aligned}$$

A *path* in G is described by $w_0 \rightarrow_{k_1} w_1 \rightarrow_{k_2} \dots \rightarrow_{k_n} w_n$, where $w_0, w_1, \dots, w_n \in V$ and $n, k_1, k_2, \dots, k_n \in \mathbb{N}$. An *access path* of a vertex w of G is a path that starts at the root of G , ends in w , and does not visit any vertex twice. Access paths need not be unique. A term graph is *root-connected* if every vertex has an access path.

Note: By a ‘term graph’ we will, from now on, always mean a root-connected term graph.

Let $G_1 = \langle V_1, lab_1, args_1, r_1 \rangle$, $G_2 = \langle V_2, lab_2, args_2, r_2 \rangle$ be term graphs over signature Σ , in the sequel.

A *bisimulation* between G_1 and G_2 is a relation $R \subseteq V_1 \times V_2$ such that the following conditions hold, for all $\langle w, w' \rangle \in R$:

$$\left. \begin{aligned} \langle r_1, r_2 \rangle &\in R && \text{(roots)} \\ lab_1(w) &= lab_2(w') && \text{(labels)} \\ \langle args_1(w), args_2(w') \rangle &\in R^* && \text{(arguments)} \end{aligned} \right\} \quad (1)$$

where the extension $R^* \subseteq V_1^* \times V_2^*$ of R to a relation between words over V_1 and words over V_2 is defined as:

$$\begin{aligned} R^* &:= \{ \langle w_1 \dots w_k, w'_1 \dots w'_k \rangle \mid \\ &\quad w_1, \dots, w_k \in V_1, w'_1, \dots, w'_k \in V_2, \\ &\quad \text{for } k \in \mathbb{N} \text{ such that } \langle w_i, w'_i \rangle \in R \text{ for all } 1 \leq i \leq k \}. \end{aligned}$$

We write $G_1 \Leftrightarrow G_2$ if there is a bisimulation between G_1 and G_2 , and we say, in this case, that G_1 and G_2 are *bisimilar*. Bisimilarity \Leftrightarrow is an equivalence relation on term graphs.

A *functional bisimulation* from G_1 to G_2 is a bisimulation that is the graph of a function from V_1 to V_2 . An alternative characterisation of this concept is that of *homomorphism* from G_1 to G_2 : a morphism from the structure G_1 to the structure G_2 , that is, a function $h : V_1 \rightarrow V_2$ such that, for all $v \in V_1$ it holds:

$$\left. \begin{aligned} h(r_1) &= r_2 && \text{(roots)} \\ lab_1(v) &= lab_2(h(v)) && \text{(labels)} \\ h^*(args_1(v)) &= args_2(h(v)) && \text{(arguments)} \end{aligned} \right\} \quad (2)$$

where h^* is the homomorphic extension $h^* : V_1^* \rightarrow V_2^*$, $v_1 \dots v_n \mapsto h(v_1) \dots h(v_n)$ of h to words over V_1 . We write $G_1 \rightrightarrows G_2$ if there is a functional bisimulation (a homomorphism) from G_1 to G_2 . An *isomorphism* between G_1 and G_2 is a bijective homomorphism $i : V_1 \rightarrow V_2$ from G_1 to G_2 . If there is an isomorphism between G_1 and G_2 , we write $G_1 \simeq G_2$, and say that G_1 and G_2 are *isomorph*.

Let $G = \langle V, lab, args, r \rangle$ be a term graph. A *bisimulation collapse* of G is a maximal element in the class $\{G' \mid G \rightrightarrows G'\}$ up to \simeq , that is, a term graph G'_0 with $G \rightrightarrows G'_0$ such that if $G'_0 \rightrightarrows G''_0$ for some term graph G''_0 , then $G'_0 \simeq G''_0$. The *canonical bisimulation collapse* $G \downarrow$ of G is defined as the root-connected part of the ‘factor term graph’ G'_R of G with respect to the largest bisimulation R between G and G (the largest ‘self-bisimulation’ on G), which is an equivalence relation on V . The *factor term graph* $G \downarrow_{\sim}$ of G with respect to an equivalence relation \sim on V is defined as $G \downarrow_{\sim} := \langle V/\sim, lab/\sim, args/\sim, [r]_{\sim} \rangle$ where V/\sim is the set of \sim -equiva-

lence classes of vertices in V , $[r]_{\sim}$ is the \sim -equivalence class of r , and lab/\sim and $args/\sim$ are the mappings on V/\sim that are induced by lab and $args$, respectively. Every two bisimulation collapses of G are isomorphic. This justifies the common abbreviation of saying that ‘the bisimulation collapse’ of G is unique up to isomorphism.

3. Unfolding Semantics of λ_{letrec} -terms

Informally, we regard λ_{letrec} -terms as being defined by the following grammar:

$$\begin{aligned} L &::= \lambda x. L && \text{(abstraction)} \\ &\mid LL && \text{(application)} \\ &\mid x && \text{(variable)} \\ &\mid \text{let } B \text{ in } L && \text{(letrec)} \\ B &::= f_1 = L, \dots, f_n = L && \text{(equations)} \\ &\quad (f_1, \dots, f_n \in \mathcal{R} \text{ all distinct}) \end{aligned}$$

Formally, we consider λ_{letrec} -terms to be defined correspondingly as higher-order terms in the formalism of Combinatory Reduction Systems (CRS) [32]. CRSs are a higher-order term rewriting framework tailor-made for formalising and manipulating expressions in higher-order languages (i.e. languages with binding constructs like λ -abstractions and let-bindings). They provide a sound basis for defining our language and for reasoning with letrec-expressions. By formalising a system of unfolding rules as a CRS we conveniently externalise issues like name capturing and α -renaming, which otherwise would have to be handled by a calculus of explicit substitution. Also, we can lean on the rewriting theory of CRSs for the proofs.

As CRS-signature we use $\Sigma_{\lambda_{\text{letrec}}} = \Sigma_{\lambda} \cup \{\text{let}_n, \text{rec-in}_n \mid n \in \mathbb{N}\}$ with $\Sigma_{\lambda} = \{\text{abs}, \text{app}\}$, where the unary symbol abs and the binary symbol app represent λ -abstraction and application, respectively; the symbols let_n of arity one, and rec-in_n of arity $n + 1$ together formalise let-expressions with n bindings. By $|L|$ we denote the size (number of symbols) of a λ_{letrec} -term L . By $\text{Ter}(\lambda_{\text{letrec}})$ we denote the set of CRS-terms over $\Sigma_{\lambda_{\text{letrec}}}$. For readability, we will rely on the informal first-order notation.

Infinite λ -terms are formalised as iCRS-terms (terms in an infinitary CRS [23]) over Σ_{λ} , forming the set $\text{Ter}(\lambda^{\infty})$. Informally, infinite λ -terms are generated co-inductively by the alternatives (*abstraction*), (*application*), and (*variable*) of the grammar above.

In order to formally define the infinite unfolding of λ_{letrec} -terms we utilise a CRS whose rewrite rules formalise unfolding steps [11]. Every λ_{letrec} -term L that represents an infinite λ -term M can be rewritten by a typically infinite rewrite sequence that converges to M in the limit. However, not every λ_{letrec} -term represents an λ^{∞} -term. For instance the λ_{letrec} -term $Q = \lambda x. \text{let } f = f \text{ in } f x$ with a meaningless let-binding for f does not unfold to a λ^{∞} -term. Therefore we introduce a constant symbol \bullet , called ‘black hole’, for expressing meaningless bindings, in order to define the unfolding operation as a total function. The unfolding semantics of Q will then be $\lambda x. \bullet x$. So we extend the signature Σ_{λ} to $\Sigma_{\lambda, \bullet}$ including \bullet , and denote the set of infinite λ -terms over Σ_{λ} by $\text{Ter}(\lambda^{\infty})$. Similarly, the rules below are defined for terms in $\text{Ter}(\lambda_{\text{letrec}, \bullet})$ based on signature $\Sigma_{\lambda_{\text{letrec}, \bullet}}$ that extends $\Sigma_{\lambda_{\text{letrec}}}$ by the blackhole constant.

Definition 3.1 (unfolding CRS for λ_{letrec} -terms). The rules

$$\begin{aligned} (\text{@}) \quad & \text{let } B \text{ in } L_0 L_1 \rightarrow (\text{let } B \text{ in } L_0) (\text{let } B \text{ in } L_1) \\ (\lambda) \quad & \text{let } B \text{ in } \lambda x. L_0 \rightarrow \lambda x. \text{let } B \text{ in } L_0 \\ (\text{let.in}) \quad & \text{let } B_0 \text{ in let } B_1 \text{ in } L \rightarrow \text{let } B_0, B_1 \text{ in } L \\ (\text{let-rec}) \quad & \text{let } B_1, f = L, B_2 \text{ in } f \rightarrow \text{let } B_1, f = L, B_2 \text{ in } L \\ (\text{gc}) \quad & \text{let } f_1 = L_1, \dots, f_n = L_n \text{ in } P \rightarrow P \\ & \quad (\text{if } f_1, \dots, f_n \text{ do not occur in } P) \end{aligned}$$

(tighten) let $B_1, f = g, B_2$ in L
 \rightarrow let $B_1[f := g], B_2[f := g]$ in $L[f := g]$
 (where g with $g \neq f$ a recursion variable in B_1 or B_2)
 (•) let $B_1, f = f, B_2$ in $L \rightarrow$ let $B_1, f = \bullet, B_2$ in L

define, in informal notation, the *unfolding CRS* for λ_{letrec} -terms with rewrite relation \rightarrow_{unf} . See Fig. 3 for the formulation of these rules in explicit CRS-notation.

Example 3.2 (Unfolding derivation of L from Ex. 1.2).

$\lambda f. \text{let } r = f r \text{ in } r \rightarrow_{\text{unf}}^{(\text{let-rec})} \lambda f. \text{let } r = f r \text{ in } f r \rightarrow_{\text{unf}}^{(\text{let-rec})} \lambda f. (\text{let } r = f r \text{ in } f) (\text{let } r = f r \text{ in } r) \rightarrow_{\text{unf}}^{(\text{gc})} \lambda f. f (\text{let } r = f r \text{ in } r) \rightarrow_{\text{unf}}^{(\text{let-rec})} \dots$

We say that a λ_{letrec} -term L *unfolds to* an $\lambda_{\bullet}^{\infty}$ -term M , or that L *expresses* M , if there is a (typically) infinite \rightarrow_{unf} -rewrite sequence from L that converges to M , symbolically $L \rightarrow_{\text{unf}}^{\infty} M$. Note that any such rewrite sequence is strongly convergent (the depth of the contracted redexes tends to infinity), because the resulting term does not contain any let-expressions.

Lemma 3.3. Every λ_{letrec} -term unfolds to precisely one $\lambda_{\bullet}^{\infty}$ -term.

Proof (Outline). Infinite normal forms of \rightarrow_{unf} are $\lambda_{\bullet}^{\infty}$ -terms since: every occurrence of a let-expression in a $\lambda_{\text{letrec}, \bullet}$ gives rise to a redex; and infinite terms over $\Sigma_{\lambda_{\text{letrec}, \bullet}}$ without let-expressions are $\lambda_{\bullet}^{\infty}$ -terms. Also, outermost-fair rewrite sequences in which the rules (tighten) and (•) are applied eagerly are (strongly) convergent.

Unique infinite normalisation of \rightarrow_{unf} follows from finitary confluence of \rightarrow_{unf} . In previous work [11] we proved confluence for the slightly simpler CRS that does not contain the final two rules, which together introduce black holes in terms with meaningless bindings. That confluence proof can be adapted by extending the argumentation to deal with the additional critical pairs. \square

Definition 3.4. The *unfolding semantics* for λ_{letrec} -terms is defined by the function $\llbracket \cdot \rrbracket_{\lambda_{\bullet}^{\infty}} : \text{Ter}(\lambda_{\text{letrec}}) \rightarrow \text{Ter}(\lambda_{\bullet}^{\infty})$, where $L \mapsto \llbracket L \rrbracket_{\lambda_{\bullet}^{\infty}} :=$ the infinite unfolding of L .

Remark 3.5 (Regular and strongly regular λ^{∞} -terms). λ^{∞} -terms that arise as infinite unfoldings of λ_{letrec} -terms form a proper subclass of those λ^{∞} -terms that have a regular term structure [11]. λ^{∞} -terms that belong to this subclass are called ‘strongly regular’, and can be characterised by means of a decomposition rewrite system, and as those that contain only finite ‘binding–capturing chains’ [11, 13].

4. Lambda higher-order term graphs

In this section we motivate the use of higher-order term graphs as a semantics for λ_{letrec} -terms; we introduce the class \mathcal{H} of ‘ λ -ho-term-graphs’ and define the semantics $\llbracket \cdot \rrbracket_{\mathcal{H}}$ for interpreting λ_{letrec} -terms as λ -ho-term-graphs. Finally, we sketch a proof of the correctness of $\llbracket \cdot \rrbracket_{\mathcal{H}}$ with respect to unfolding equivalence (the property (P1)).

We start out from a natural interpretation of λ_{letrec} -terms as first-order term graphs: occurrences of abstraction variables are resolved as edges pointing to the corresponding abstraction; occurrences of recursion variables as edges to the subgraph belonging to the respective binding. We therefore consider term graphs over the signature $\Sigma_{\bullet}^{\lambda} = \{\textcircled{\lambda}, \lambda, 0, \bullet\}$ with arities $\text{ar}(\textcircled{\lambda}) = 2$, $\text{ar}(\lambda) = 1$, $\text{ar}(0) = 1$, and $\text{ar}(\bullet) = 0$. These function symbols represent applications, λ -abstractions, abstraction variables, and black holes.

We will later define a subclass of these term graphs that excludes meaningless graphs. In line with the choice to regard all terms as higher-order terms (thus modulo α -conversion), we consider a nameless graph representation, so that α -equivalence of two terms can be recognised as their graph interpretations being isomorphic.

For a term graph G over $\Sigma_{\bullet}^{\lambda}$ with set V of vertices we will henceforth denote by $V(\textcircled{\lambda})$, $V(\lambda)$, $V(0)$, and $V(\bullet)$ the sets of *application vertices*, *abstraction vertices*, *variable vertices*, and *blackhole vertices*, that is, those with label $\textcircled{\lambda}$, λ , 0 , \bullet , respectively.

Example 4.1 (Natural first-order interpretation). The λ_{letrec} -terms L and P in Ex. 1.2 can be represented as the term graphs in Fig. 2.

These two graphs are bisimilar, which suggests that L and P are unfolding equivalent. Moreover, there is a functional bisimulation from the larger term graph to the smaller one, indicating that L expresses more sharing than P , or in other words: L is more compact. Also, there is no smaller term graph that is bisimilar to L and P . We conclude that L is a maximally shared form of P .

However, this translation is incorrect in the sense that bisimilarity does not in general guarantee unfolding equivalence, the desired property (P1). This is witnessed by the following counterexample.


Example 4.2 (Incorrectness of the natural first-order interpretation).

$L_1 = \text{let } f = \lambda x. (\lambda y. f y) x \text{ in } f$

$L = \text{let } f = \lambda x. f x \text{ in } f$

$L_2 = \text{let } f = \lambda x. (\lambda y. f x) x \text{ in } f$

While $\llbracket L_1 \rrbracket_{\lambda_{\bullet}^{\infty}} = \llbracket L \rrbracket_{\lambda_{\bullet}^{\infty}}$ and $\llbracket L \rrbracket_{\lambda_{\bullet}^{\infty}} \neq \llbracket L_2 \rrbracket_{\lambda_{\bullet}^{\infty}}$, all of their term graphs G_1, G, G_2 are bisimilar (please ignore the shading for now):



Consequently this interpretation lacks the necessary structure for correctly modelling compactification via bisimulation collapse.

We therefore impose additional structure on the term graphs. This is indicated by the shading in the picture above, and in the graphs throughout this paper. A shaded area depicts the *scope* of an abstraction: it comprises all positions between the abstraction and its bound variable occurrences as well as the scope of any abstraction on these positions. By this stipulation, scopes are properly nested.

Now note that the functional bisimulation on the right in the picture in Ex. 4.2 does *not* respect the scopes: The scope of the topmost abstraction vertex in the term graph G_2 interpreting L_2 contains another λ -abstraction; hence the image of this scope under the functional bisimulation cannot fit into, and is not contained in, the single scope in the term graph G of L . Also, the trivial scope of the vacuous abstraction in G_2 is not mapped to a scope in G . Thus the natural first-order interpretation is incorrect, in the sense that functional bisimulation does not preserve scopes on the first-order term graphs that are interpretations of λ_{letrec} -terms.


To prevent that interpretations of not unfolding-equivalent terms like L_1 and L_2 in Ex. 4.2 become bisimilar, we enrich first-order term graphs by a formal concept of scope. More precisely, *abstraction prefixes* are added as vertex labels. They also serve the purpose of defining the subclass of meaningful term graphs over $\Sigma_{\bullet}^{\lambda}$ that sensibly represent cyclic λ -terms. In the enriched term graphs, each vertex v is annotated with a label $P(v)$, the *abstraction prefix* of v , which is a list of vertex names that identifies the abstraction vertices in whose scope v resides. Alternatively scopes can be represented by a scope function (as in [4]) that assigns to every abstraction vertex the set of vertices in its scope. In the article [12] we show that higher-order term graphs with scope functions correspond bijectively to those with abstraction prefix functions.

(@)	$\text{let}_n([\vec{f}] \text{rec-in}_n(X_1(\vec{f}), \dots, X_n(\vec{f}), \text{app}(Z_0(\vec{f}), Z_1(\vec{f}))))$ $\rightarrow \text{app}((\text{let}_n([\vec{f}] \text{rec-in}_n(\dots, X_n(\vec{f}), Z_0(\vec{f}))), (\text{let}_n([\vec{f}] \text{rec-in}_n(\dots, X_n(\vec{f}), Z_1(\vec{f}))))))$
(λ)	$\text{let}_n([\vec{f}] \text{rec-in}_n(X_1(\vec{f}), \dots, X_n(\vec{f}), \text{abs}([x] Z(\vec{f}, x))))$ $\rightarrow \text{abs}([x] \text{let}_n([\vec{f}] \text{rec-in}_n(X_1(\vec{f}), \dots, X_n(\vec{f}), Z(\vec{f}, x))))$
(let_in)	$\text{let}_n([\vec{f}] \text{rec-in}_n(\vec{X}(\vec{f}), \text{let}_m([\vec{g}] \text{rec-in}_m(\vec{Y}(\vec{f}, \vec{g})), Z(\vec{f}, \vec{g})))$ $\rightarrow \text{let}_{n+m}([\vec{f}\vec{g}] \text{rec-in}_{n+m}(\vec{X}(\vec{f}), \vec{Y}(\vec{f}, \vec{g}), Z(\vec{f}, \vec{g})))$
(let-rec)	$\text{let}_n([\vec{f}] \text{rec-in}_n(X_1(\vec{f}), \dots, X_n(\vec{f}), f_i)) \rightarrow \text{let}_n([\vec{f}] \text{rec-in}_n(X_1(\vec{f}), \dots, X_n(\vec{f}), X_i(\vec{f})))$
(gc)	$\text{let}_n([\vec{f}] \text{rec-in}_n(X_1(\vec{f}), \dots, X_n(\vec{f}), Z)) \rightarrow Z$
(tighten)	$\text{let}_n([\vec{f}] \text{rec-in}_n(X_1(\vec{f}), \dots, X_{i-1}(\vec{f}), f_j, X_{i+1}(\vec{f}), \dots, X_n(\vec{f}), Z(\vec{f})))$ $\rightarrow \text{let}_{n-1}([\vec{g}] \text{rec-in}_{n-1}(X_1(\vec{g}'), \dots, X_{i-1}(\vec{g}'), X_{i+1}(\vec{g}'), \dots, X_n(\vec{g}'), Z(\vec{g}')))$ <p style="text-align: center;">where: $i, j \in \{1, \dots, n\}$ with $i \neq j$, and $\vec{g}' = \langle g_1, \dots, g_{i-1}, g_j, g_{i+1}, \dots, g_n \rangle$</p>
(•)	$\text{let}_n([\vec{f}] \text{rec-in}_n(X_1(\vec{f}), \dots, X_{i-1}(\vec{f}), f_i, X_{i+1}(\vec{f}), \dots, X_n(\vec{f}), f_i))$ $\rightarrow \text{let}_n([\vec{f}] \text{rec-in}_n(X_1(\vec{f}), \dots, X_{i-1}(\vec{f}), \bullet, X_{i+1}(\vec{f}), \dots, X_n(\vec{f}), f_i))$

Figure 3. The rules for unfolding λ_{letrec} -terms in explicit CRS-notation.

Abstraction prefixes can be determined by traversing over the graph and recording every binding encountered. When passing an abstraction vertex v while descending into the subgraph representing the body of the abstraction, one enters or opens the scope of v . This is recorded by appending v to the abstraction prefix of v 's successor. v is removed from the prefix at positions under which the abstraction variable is no longer used, but not before any other variable that was added to the prefix in the meantime has itself been removed. In other words, the abstraction prefix behaves like a stack. We call term graphs for representing λ_{letrec} -terms that are equipped with abstraction-prefixes ‘λ-higher-order term graphs’ (λ-ho-term-graphs).

Example 4.3 (The λ-ho-term-graphs of the terms in Ex. 4.2).



The superscripts of abstraction vertices indicate their names. The abstraction prefix of a vertex is annotated to its top left. Note that abstraction vertices themselves are not included in their own prefix.

We define λ-ho-term-graphs as term graphs over Σ_\bullet^λ together with an abstraction-prefix function that assigns to each vertex an abstraction prefix. It has to respect certain correctness conditions restricting the λ-ho-term-graphs to exclude meaningless term graphs.

Definition 4.4 (correct abstraction-prefix function for term graphs over Σ_\bullet^λ). Let $G = \langle V, \text{lab}, \text{args}, r \rangle$ be a Σ_\bullet^λ -term-graph.

An *abstraction-prefix function* for G is a function $P : V \rightarrow V^*$ from vertices of G to words of vertices. Such a function is called

correct if for all $w, w_0, w_1 \in V$ and $k \in \{0, 1\}$ it holds:

$$\begin{aligned}
 P(r) &= \epsilon & (\text{root}) \\
 P(\bullet) &= \epsilon & (\text{black hole}) \\
 w \in V(\lambda) \wedge w \twoheadrightarrow_0 w_0 &\Rightarrow P(w_0) \leq P(w)w & (\lambda) \\
 w \in V(@) \wedge w \twoheadrightarrow_k w_k &\Rightarrow P(w_k) \leq P(w) & (@) \\
 w \in V(0) \wedge w \twoheadrightarrow_0 w_0 &\Rightarrow \begin{cases} w_0 \in V(\lambda) \\ \wedge P(w_0)w_0 = P(w) \end{cases} & (0)
 \end{aligned}$$

Here and later we denote by \leq the ‘is-prefix-of’ relation.

Definition 4.5 (λ-ho-term-graph). A λ-ho-term-graph over Σ_\bullet^λ is a five-tuple $\mathcal{G} = \langle V, \text{lab}, \text{args}, r, P \rangle$ where $G_{\mathcal{G}} = \langle V, \text{lab}, \text{args}, r \rangle$ is a term graph over Σ_\bullet^λ , called the term graph *underlying* \mathcal{G} , and P is a correct abstraction-prefix function for $G_{\mathcal{G}}$. The class of λ-ho-term-graphs over Σ_\bullet^λ is denoted by \mathcal{H} .

Definition 4.6 (homomorphism, bisimulation for λ-ho-term-graphs). Let $\mathcal{G}_1 = \langle V_1, \text{lab}_1, \text{args}_1, r_1, P_1 \rangle$ and $\mathcal{G}_2 = \langle V_2, \text{lab}_2, \text{args}_2, r_2, P_2 \rangle$ be λ-ho-term-graphs over Σ_\bullet^λ .

A *bisimulation* between \mathcal{G}_1 and \mathcal{G}_2 is a relation $R \subseteq V_1 \times V_2$ such that for all $\langle w, w' \rangle \in R$ the conditions (1), and additionally: A *bisimulation* between \mathcal{G}_1 and \mathcal{G}_2 is a relation $R \subseteq V_1 \times V_2$ that is a bisimulation between the term graphs $G_{\mathcal{G}_1}$ and $G_{\mathcal{G}_2}$ underlying \mathcal{G}_1 and \mathcal{G}_2 , respectively, and for which also the following condition:

$$\langle P_1(w), P_2(w') \rangle \in R^* \quad (\text{abstraction-prefix functions}) \quad (3)$$

(for R^* see p. 4 below (1)) is satisfied for all $w \in V_1$ and all $w' \in V_2$. If there is a such bisimulation, then \mathcal{G}_1 and \mathcal{G}_2 are *bisimilar*, denoted by $\mathcal{G}_1 \rightleftharpoons \mathcal{G}_2$.

A *homomorphism* (a *functional bisimulation*) from \mathcal{G}_1 to \mathcal{G}_2 is a morphism from the structure \mathcal{G}_1 to the structure \mathcal{G}_2 , or more explicitly, it is a homomorphism $h : V_1 \rightarrow V_2$ from $G_{\mathcal{G}_1}$ to $G_{\mathcal{G}_2}$ that additionally satisfies, for all $w \in V_1$, the following condition:

$$\bar{h}(P_1(w)) = P_2(h(w)) \quad (\text{abstraction-prefix functions}) \quad (4)$$

for all $w \in V_1$, where \bar{h} is the homomorphic extension of h to words over V_1 . We write $\mathcal{G}_1 \Rightarrow \mathcal{G}_2$ if there is a homomorphism between \mathcal{G}_1 and \mathcal{G}_2 .




Figure 4. Translation rules \mathcal{R} for interpreting λ_{letrec} -terms as λ -ho-term-graphs. See Section 4.1 for explanations.

4.1 Interpretation of λ_{letrec} -terms as λ -ho-term-graphs

In order to interpret a λ_{letrec} -term L as λ -ho-term-graph, the translation rules \mathcal{R} from Fig. 4 are applied to a ‘translation box’ $(\ast[\]L)$. It contains L furnished with a prefix consisting of a dummy variable \ast , and an empty set $[\]$ of binding equations. The translation process proceeds by induction on the syntactical structure of the prefixed λ_{letrec} -expression’s body. Ultimately, a term graph G over $\Sigma_{\bullet}^{\lambda}$ is produced, together with a correct abstraction-prefix function for G .

For reading the rules \mathcal{R} in Fig. 4 correctly, observe the details as described here below. For illustration of their application, please refer to Appendix A where several λ_{letrec} -terms are translated into λ -ho-term-graphs.

- A translation box $(\ast[\]L)$ contains a prefixed, partially decomposed λ_{letrec} -term L . The prefix contains a vector \vec{p} of annotated λ -abstractions that have already been translated and whose scope typically extends into L . Every prefix abstraction is annotated with a set of binding equations that are defined at its level. There is special dummy variable denoted by \ast at the left of the prefix that carries top-level function bindings, i.e. binding equations that are not defined under any enclosing λ -abstraction. The λ -rule strips off an abstraction from the body of the expression, and pushes the abstraction variable into the prefix, which initially contains an empty set of function bindings.
- Names of abstraction vertices are indicated to the right, and abstraction-prefixes to the left of the created vertices. In order to refer to the vertices in the prefix we use the following notation: $vs(\vec{p}) = v_1 \dots v_n$ if $\vec{p} = \ast[B_0] x_1^{v_1}[B_1] \dots x_n^{v_n}[B_n]$.
- Vertices drawn with dashed lines have been created in earlier translation steps, and in the current step are referenced by edges in the current step.
- In the S-rule, which takes care of closing scopes, $FV(L)$ stands for the set of free variables in L .
- The let-rule for translating let-expressions creates a box for the in-part as well as for each binding equation. The translation of each of the bindings starts with an *indirection* vertex. These vertices guarantee the well-definedness of the process when it translates meaningless bindings such as $f = f$, or $g = h$, $h = g$, which would otherwise give rise to loops without vertices. The let-rule pushes the function bindings into the abstraction prefix, associating each function binding with one of the variables in the

abstraction prefix. There is some freedom as to which variable a function binding is assigned to. This freedom is limited by scoping conditions that ensure that the prefixed term is a valid CRS-term: function bindings may only depend on variables and functions that occur further to the left in the prefix. The chosen association also directly determines the prefix lengths used in the translation boxes for the function bindings.

- Indirection vertices are eliminated by an erasure process at the end: Every indirection vertex that does not point to itself is removed, redirecting all incoming edges to the successor vertex. Finally every loop on a single indirection vertex is replaced by a *black hole* vertex that represents a meaningless binding. Abstraction prefixes for such black holes are defined to be empty.




Figure 5. Translation of $\lambda a. \lambda b. \text{let } f = a \text{ in } a a (f a) b$ with equal (left) and with minimal prefix lengths (right) in the let-rule.

Definition 4.7. We say that a term graph G over $\Sigma_{\bullet}^{\lambda}$ and an abstraction-prefix function P is \mathcal{R} -generated from a λ_{letrec} -term L if G and P are obtained by applying the rules \mathcal{R} from Fig. 4 to $(\ast[\]L)$.

Remark 4.8 (Inference rule formulation of \mathcal{R}). See also Fig. 6 for inference rules that correspond to the deconstruction of prefixed terms in \mathcal{R} .

Proposition 4.9. Let L be a λ_{letrec} -term. Suppose that a term graph G over $\Sigma_{\bullet}^{\lambda}$, and an abstraction-prefix function P are \mathcal{R} -generated from L . Then P is a correct abstraction-prefix function for G , and consequently, G and P together form a λ -ho-term-graph in \mathcal{H} .

$$\begin{array}{c}
\frac{(\bar{p} \ x^v[]) L}{(\bar{p}) \lambda x. L} \lambda \quad \frac{(\bar{p}) L_0 \quad (\bar{p}) L_1}{(\bar{p}) L_0 L_1} @ \quad \frac{}{(\bar{p} \ x^v[\dots, f^w = L, \dots]) f} \text{rec} \\
\\
\frac{(x_0^{v_0}[B_0] \ \dots \ x_n^{v_n}[B_n]) x_n}{0} 0 \quad \frac{(\bar{p}) L}{(\bar{p} \ x^v[f_1^{v_1} = L_1, \dots, f_n^{v_n} = L_n]) L} S \text{ (if } x \notin FV(L) \text{ and } f_i \notin FV(L)) \\
\\
\frac{(x_0^{v_0}[B'_0] \ \dots \ x_{l_1}^{v_{l_1}}[B'_{l_1}]) L_1 \quad \dots \quad (x_0^{v_0}[B'_0] \ \dots \ x_{l_k}^{v_{l_k}}[B'_{l_k}]) L_k \quad (x_0^{v_0}[B'_0] \ \dots \ x_n^{v_n}[B'_n]) L_0}{(x_0^{v_0}[B_0] \ \dots \ x_n^{v_n}[B_n]) \text{let } f_1 = L_1, \dots, f_k = L_k \text{ in } L_0} \text{let (where } B'_0, \dots, B'_n \text{ and } l_1, \dots, l_k \text{ as in rule let in Fig. 4)}
\end{array}$$

Figure 6. Alternative formulation as inference rules of the translation rules in Fig. 4 for the interpretation of λ_{letrec} -terms as λ -ho-term-graphs.




Figure 7. Translation of $\lambda a. (\lambda b. \lambda c. a c) (\lambda d. a d)$ with eager scope-closure (left), and with lazy scope-closure (right). While in the left term graph four vertices can be shared, with as result the translation of the term $\lambda a. \text{let } f = \lambda c. a c \text{ in } (\lambda b. f) f$, in the right term graph only a single variable occurrence can be shared.

There are two sources of non-determinism in this translation: The S-rule for shortening prefixes can be applicable at the same time as other rules. And the let-rule does not fix the lengths l_1, \dots, l_k of the abstraction prefixes for the translations of the binding equations, but admits various choices of prefixes that are shorter than the prefix of the left-hand side. Neither kind of non-determinism affects the term graph that is produced, but in general several abstraction-prefix functions, and thus different λ -ho-term-graphs, can be obtained.

4.2 Interpretation as eager-scope λ -ho-term-graphs

Of the different translations of a λ_{letrec} -term into λ -ho-term-graphs we are most interested in the one with the shortest possible abstraction prefixes. We say that such a term graph has ‘eager scope-closure’, or that it is ‘eager-scope’. The reason for this choice is illustrated in Fig. 7: eager-scope closure allows for more sharing.

Definition 4.10 (eager scope). Let $\mathcal{G} = \langle V, \text{lab}, \text{args}, r, P \rangle$ be a λ -ho-term-graph. \mathcal{G} is called *eager-scope* if for every $w \in V$ with $P(w) = pv$ for $p \in V^*$ and $v \in V$, there is a path $w = w_0 \rightsquigarrow w_1 \rightsquigarrow \dots \rightsquigarrow w_m \rightsquigarrow_0 v$ in \mathcal{G} from w to v with $P(w) \leq P(w_i)$ for all $i \in \{1, \dots, m\}$, and (this follows) $w_m \in V(0)$ and $v \in V(\lambda)$.

Hence if a λ -ho-term-graph is not eager-scope, then it contains a vertex w with abstraction-prefix $v_1 \dots v_n$ from which v_n is only reachable, if at all, by leaving the scope of v_n . It can be shown that in this case another abstraction-prefix function with shorter prefixes exists, and in which v_n has been removed from the prefix of w .

Proposition 4.11 (eager-scope = minimal scope; uniqueness of eager-scope λ -ho-term-graphs). Let $\mathcal{G}_i = \langle V, \text{lab}, \text{args}, r, P_i \rangle$ for $i \in \{1, 2\}$ be λ -ho-term-graphs with the same underlying term graph.

If \mathcal{G}_1 is eager-scope, then $|P_1(w)| \leq |P_2(w)|$ for all $w \in V$. If, in addition, also \mathcal{G}_2 is eager-scope, then $P_1 = P_2$. Hence eager-scope λ -ho-term-graphs over the same underlying term graph are unique.

Also, we will call a translation process ‘eager-scope’ if it resolves the non-determinism in \mathcal{R} in such a way that it always yields eager-scope λ -ho-term-graphs. In order to obtain an eager-scope translation we have to consider the following aspects.

Garbage removal. In the presence of *garbage*, unused function bindings, a translation cannot be eager-scope. Consider the term $\lambda x. \lambda y. \text{let } f = x \text{ in } y$. The expendable binding $f = x$ prevents the application of the S-rule, and hence the closure of the scope of λx , directly below λx . Therefore we henceforth assume that *all unused function bindings are removed* prior to applying the rules \mathcal{R} . A λ_{letrec} -term without garbage will be called *garbage-free*.

Short enough prefix lengths in the let-rule. For obtaining an eager-scope translation, we will usually stipulate that the S-rule is applied eagerly, i.e. it is given precedence over the other rules. This is clearly necessary for keeping the abstraction prefixes minimal. But how do we choose the prefix lengths l_1, \dots, l_k in the let-rule? The prefix lengths l_i determine at which position a binding $f_i = L_i$ is inserted into the abstraction prefixes. Therefore l_i may not be chosen too short; otherwise a function f depending on a function g may end up to the right of g , and hence may be removed from the prefix by the S-rule prematurely, preventing completion of the translation. Yet simply choosing $l_i = n$ may prevent scopes from being minimal. For example, when translating the term $\lambda a. \lambda b. \text{let } f = a \text{ in } a a (f a) b$, it is crucial to allow shorter prefixes for the binding than for the in-part. As shown in Fig. 5 the graph on the left does not have eager scope-closure even if the S-rule is applied eagerly. Consequently the opportunity for sharing the lower application vertices is lost.

Required variable analysis. For choosing the prefixes in the let-rule correctly, the translation process must know for each function binding which λ -variables are ‘required’ on the right-hand side of its definition. For this we use an analysis obtaining the required variables for positions in a λ_{letrec} -term as employed by algorithms for lambda-lifting [7, 22]. The term ‘required variables’ was coined by Morazán and Schultz [25]. A λ -variable x is called *required at a position p* in a λ_{letrec} -term L if x is bound by an abstraction above p , and has a free occurrence in the complete unfolding of L below p (also recursion variables from above p are unfolded). The required variables at position p in L can be computed as those λ -variables with free occurrences that are reachable from p by a downwards traversal with the stipulations: on encountering a let-binding the in-part is entered; when encountering a recursion variable the traversal continues at the right-hand side of the corresponding function binding (even if it is defined above p).

With the result of the required variable analysis at hand, we now define properties of the translation process that can guarantee that the resulting λ -ho-term-graph is eager-scope.

Definition 4.12 (eager-scope and minimal-prefix generated). Let L be a λ_{letrec} -term, and let \mathcal{G} be a λ -ho-term-graph.

We say that \mathcal{G} is *eager-scope* \mathcal{R} -generated from L if \mathcal{G} is \mathcal{R} -generated from L by a translation process with the following property: for every translation box reached during the process with label $(\bar{p} x^v[B]) P$, where P is a subterm of L at position q , it holds that if x is not a required variable at q in L , then in the next translation step performed to this box either one of the rules f or let is applied, or the prefix is shortened by the S -rule.

We say that \mathcal{G} is \mathcal{R} -generated with *minimal prefixes* from L if \mathcal{G} is \mathcal{R} -generated from L by a translation process in which minimal prefix lengths are achieved by giving applications of the S -rule precedence over applications of all other rules, and by always choosing prefixes minimally in applications of the let -rule.

Proposition 4.13. Let \mathcal{G} be a λ -ho-term-graph that is \mathcal{R} -generated from a garbage-free λ_{letrec} -term L . The following statements hold:

- (i) If \mathcal{G} is eager-scope \mathcal{R} -generated from L , then \mathcal{G} is eager-scope.
- (ii) If \mathcal{G} is \mathcal{R} -generated with minimal prefixes from L , then \mathcal{G} is eager-scope \mathcal{R} -generated from L , hence by (i) \mathcal{G} is eager-scope.

Proposition 4.14. For every λ_{letrec} -term L , $\llbracket L \rrbracket_{\mathcal{H}}$ is eager-scope.

4.3 Correctness of $\llbracket \cdot \rrbracket_{\mathcal{H}}$ with respect to unfolding semantics

In preparation of establishing the desired property (P1) in Sect. 5, we formulate, and outline the proof of, the fact that the semantics $\llbracket \cdot \rrbracket_{\mathcal{H}}$ is correct with respect to the unfolding semantics on λ_{letrec} -terms.

Theorem 4.15. $\llbracket L_1 \rrbracket_{\lambda^\infty} = \llbracket L_2 \rrbracket_{\lambda^\infty}$ if and only if $\llbracket L_1 \rrbracket_{\mathcal{H}} \rightleftharpoons \llbracket L_2 \rrbracket_{\mathcal{H}}$, for all λ_{letrec} -terms L_1 and L_2 .

Sketch of Proof. Central for the proof are λ -ho-term-graphs that have tree form and only contain variable backlinks, but no recursive backlinks. They form the class $\mathcal{H}_T \subseteq \mathcal{H}$. Every $\mathcal{G} \in \mathcal{H}$ has a unique ‘tree unfolding’ $\text{Tree}(\mathcal{G}) \in \mathcal{H}_T$. We make use of the following statements. For all $L, L_1, L_2 \in \text{Ter}(\lambda_{\text{letrec}})$, $M, M_1, M_2 \in \text{Ter}(\lambda^\infty)$, $\mathcal{G}, \mathcal{G}_1, \mathcal{G}_2 \in \mathcal{H}$, and $\mathcal{T}r, \mathcal{T}r_1, \mathcal{T}r_2 \in \mathcal{H}_T$ it can be shown that:

$$L_1 \rightarrow_{\text{unf}} L_2 \Rightarrow \llbracket L_1 \rrbracket_{\mathcal{H}} \preceq \llbracket L_2 \rrbracket_{\mathcal{H}} \quad (5)$$

$$L \twoheadrightarrow_{\text{unf}} M \text{ (hence } \llbracket L \rrbracket_{\lambda^\infty} = M) \Rightarrow \llbracket L \rrbracket_{\mathcal{H}} \preceq \llbracket M \rrbracket_{\mathcal{H}} \quad (6)$$

$$\llbracket M \rrbracket_{\mathcal{H}} \in \mathcal{H}_T \quad (7)$$

$$\llbracket M_1 \rrbracket_{\mathcal{H}} \simeq \llbracket M_2 \rrbracket_{\mathcal{H}} \Rightarrow M_1 = M_2 \quad (8)$$

$$\mathcal{G} \preceq \text{Tree}(\mathcal{G}) \quad (9)$$

$$\mathcal{T}r_1 \preceq \mathcal{T}r_2 \Rightarrow \mathcal{T}r_1 \simeq \mathcal{T}r_2 \quad (10)$$

$$\mathcal{G}_1 \preceq \mathcal{G}_2 \Rightarrow \text{Tree}(\mathcal{G}_1) \simeq \text{Tree}(\mathcal{G}_2) \quad (11)$$

Hereby (5) is used for proving (6), and (9) with (10) for (11). Now for proving the theorem, let L_1 and L_2 be arbitrary λ_{letrec} -terms.

“ \Rightarrow ”: Suppose $\llbracket L_1 \rrbracket_{\lambda^\infty} = \llbracket L_2 \rrbracket_{\lambda^\infty}$. Let M be the infinite unfolding of L_1 and L_2 , i.e., $\llbracket L_1 \rrbracket_{\mathcal{H}} = M = \llbracket L_2 \rrbracket_{\mathcal{H}}$. Then by (6) it follows $\llbracket L_1 \rrbracket_{\mathcal{H}} \preceq \llbracket M \rrbracket_{\mathcal{H}} \Rightarrow \llbracket L_2 \rrbracket_{\mathcal{H}}$, and hence $\llbracket L_1 \rrbracket_{\mathcal{H}} \rightleftharpoons \llbracket L_2 \rrbracket_{\mathcal{H}}$.

“ \Leftarrow ”: Suppose $\llbracket L_1 \rrbracket_{\mathcal{H}} \rightleftharpoons \llbracket L_2 \rrbracket_{\mathcal{H}}$. Then by (11) it follows that $\text{Tree}(\llbracket L_1 \rrbracket_{\mathcal{H}}) \simeq \text{Tree}(\llbracket L_2 \rrbracket_{\mathcal{H}})$. Let $M_1, M_2 \in \text{Ter}(\lambda^\infty)$ be the infinite unfoldings of L_1 and L_2 , i.e. $M_1 = \llbracket L_1 \rrbracket_{\lambda^\infty}$, and $M_2 = \llbracket L_2 \rrbracket_{\lambda^\infty}$. Then (6) together with the assumption entails $\llbracket M_1 \rrbracket_{\mathcal{H}} \preceq \llbracket M_2 \rrbracket_{\mathcal{H}}$. Since $\llbracket M_1 \rrbracket_{\mathcal{H}}, \llbracket M_2 \rrbracket_{\mathcal{H}} \in \mathcal{H}_T$ by (7), it follows by (10) that $\llbracket M_1 \rrbracket_{\mathcal{H}} \simeq \llbracket M_2 \rrbracket_{\mathcal{H}}$. Finally, by using (8) we get $M_1 = M_2$, and hence $\llbracket L_1 \rrbracket_{\lambda^\infty} = M_1 = M_2 = \llbracket L_2 \rrbracket_{\lambda^\infty}$. \square


5. Lambda term graphs

While modelling sharing expressed by λ_{letrec} -terms through λ -ho-term-graphs facilitates comparisons via bisimilarity, it is not im-

mediately clear how the compactification of λ -ho-term-graphs via the bisimulation collapse \Downarrow for λ -ho-term-graphs (which has to respect scopes in the form of the abstraction-prefix functions) can be computed efficiently. We therefore develop an implementation as first-order term graphs, for which standard methods are available.

Due to Ex. 4.2, the scoping information cannot just be discarded, as functional bisimilarity on the underlying term graphs does not faithfully implement functional bisimilarity on λ -ho-term-graphs. Therefore the scoping information has to be incorporated in the first-order interpretation in some way. We accomplish this by extending Σ_{\bullet}^λ with S -vertices, scope delimiters, that signify the end of scopes. When translating a λ -ho-term-graph into a first-order term graph, S -vertices are placed along those edges in the underlying term graph at which the abstraction prefix decreases in the λ -ho-term-graph.


Example 5.1 (Adding S -vertices). Consider the terms in Ex. 4.2 and their λ -ho-term-graphs in Ex. 4.3. In the first-order interpretation below, the shading is just for illustration purposes; it is *not* part of the structure, and does *not directly* impair functional bisimulation.



The addition of scope delimiters resolves the problem of Ex.4.2. They adequately represent the scoping information.

As for λ -ho-term-graphs, we will define correctness conditions by means of an abstraction-prefix function. However, the current approach with unary delimiter vertices leads to a problem.

Example 5.2 (S -backlinks). The term graph with scope delimiters on the left admits a functional bisimulation that fuses two S -vertices that close different scopes. We cannot hope to find a unique abstraction prefix for the resulting fused S -vertex. This is remedied on the right by using a variant representation that requires backlinks from each S -vertex to the abstraction vertex whose scope it closes. Then S -vertices can only be fused if the corresponding abstractions have already been merged. Hence in the presence of S -backlinks, as in the right illustration below, only the variable vertex can be shared.



Therefore we consider term graphs over the extension $\Sigma_{S, \bullet}^\lambda$ of Σ_{\bullet}^λ with a symbol S of arity 2; one edge targets the successor vertex, the other is a backlink. We give correctness conditions, similar as for λ -ho-term-graphs, and define the arising class of ‘ λ -term-graphs’.

Definition 5.3 (correct abstraction-prefix function for term graphs over $\Sigma_{S, \bullet}^\lambda$). Let $G = \langle V, \text{lab}, \text{args}, r \rangle$ be a $\Sigma_{S, \bullet}^\lambda$ -term-graph.

An *abstraction-prefix function* $P : V \rightarrow V^*$ on G is called *correct* if for all $w, w_0, w_1 \in V$ and $k \in \{0, 1\}$ it holds:

$$\begin{aligned}
P(r) &= \epsilon & (\text{root}) \\
P(\bullet) &= \epsilon & (\text{black hole}) \\
w \in V(\lambda) \wedge w \rightarrow_0 w_0 &\Rightarrow P(w_0) = P(w)w & (\lambda) \\
w \in V(@) \wedge w \rightarrow_k w_k &\Rightarrow P(w_k) = P(w) & (@) \\
w \in V(0) \wedge w \rightarrow_0 w_0 &\Rightarrow \begin{cases} w_0 \in V(\lambda) \\ \wedge P(w_0)w_0 = P(w) \end{cases} & (0)_1 \\
w \in V(S) \wedge w \rightarrow_0 w_0 &\Rightarrow \begin{cases} P(w_0)v = P(w) \\ \text{for some } v \in V \end{cases} & (S)_1 \\
w \in V(S) \wedge w \rightarrow_1 w_1 &\Rightarrow \begin{cases} w_1 \in V(\lambda) \\ \wedge P(w_1)w_1 = P(w) \end{cases} & (S)_2
\end{aligned}$$

While in λ -ho-term-graphs the abstraction prefix can shrink by several vertices along an edge (cf. Def. 4.4), here the situation is strictly regulated: the prefix can only shrink by one variable, and only along the outgoing edge of a delimiter vertex.

Proposition 5.4 (uniqueness of the abstraction prefix function). Let G be a term graph over the signature $\Sigma_{S, \bullet}^\lambda$. If P_1 and P_2 are correct abstraction prefix functions of G , then $P_1 = P_2$.

Definition 5.5 (λ -term-graph). A λ -term-graph is a term graph $G = \langle V, \text{lab}, \text{args}, r \rangle$ over $\Sigma_{S, \bullet}^\lambda$ that has a correct abstraction-prefix function (which is not a part of G). The class of λ -term-graphs is \mathcal{T} .

Definition 5.6 (eager scope). A λ -term-graph G is called *eager-scope* if together with its abstraction-prefix function it meets the condition in Def. 4.10. \mathcal{T}_{eag} denotes the class of eager-scope graphs.

5.1 Correspondence between λ -ho- and λ -term-graphs

The correspondences between λ -ho-term-graphs and λ -term-graphs:

$$\mathcal{HT} : \mathcal{H} \rightarrow \mathcal{T} \quad \mathcal{TH} : \mathcal{T} \rightarrow \mathcal{H}$$

are defined as follows: For obtaining $\mathcal{HT}(\mathcal{G})$ for a $\mathcal{G} \in \mathcal{H}$, insert scope-delimiters wherever the prefix decreases, as illustrated in Fig. 8. For obtaining $\mathcal{TH}(G)$ for a $G \in \mathcal{T}$, retain the abstraction-prefix function, and remove every delimiter vertex from G , thereby connecting its incoming edge with its outgoing edge. For formal definitions and well-definedness of \mathcal{TH} and \mathcal{HT} , see [12].




Figure 8. Left: definition of \mathcal{HT} by inserting S -vertices, between edge-connected vertices of a λ -ho-term-graph. Right: interpretation $\mathcal{HT}(\mathcal{G})$ of the eager-scope λ -ho-term-graph \mathcal{G} in Fig. 7.

Note that a λ -ho-term-graph may have multiple corresponding λ -term-graphs that differ only with respect to their ‘degree’ of S -sharing (the extent to which S -vertices occur shared). \mathcal{HT} maps to a λ -term-graph with no sharing of S -vertices at all.

The proposition below guarantees the usefulness of the translation \mathcal{HT} for implementing functional bisimulation on λ -ho-term-graphs. In particular, this is due to items (iii) and (iv). As formulated by item (i), \mathcal{TH} is a retraction of \mathcal{HT} (and \mathcal{HT} a section of \mathcal{TH}). The converse is not the case, yet it holds up to S -sharing by item (ii). For the proof, we refer to our article [12].

Proposition 5.7 (correspondence with λ -ho-term-graphs).

- (i) $\mathcal{TH} \circ \mathcal{HT} = \text{id}_{\mathcal{H}}$.
- (ii) $(\mathcal{HT} \circ \mathcal{TH})(G) \Rightarrow^S G$ holds for all $G \in \mathcal{T}$.
- (iii) \mathcal{TH} and \mathcal{HT} preserve and reflect functional bisimulation \Rightarrow and bisimulation \Leftrightarrow on \mathcal{H} and \mathcal{T} .
- (iv) \mathcal{TH} and \mathcal{HT} preserve and reflect the property eager-scope.
- (v) \mathcal{T} is closed under \Rightarrow^S , \Leftarrow^S , and \Leftrightarrow^S .
- (vi) \mathcal{HT} and \mathcal{TH} induce isomorphisms between \mathcal{H} and $\mathcal{T}/\Leftrightarrow^S$.

5.2 Closedness of \mathcal{T} under functional bisimulation

While preservation of \Rightarrow by \mathcal{HT} is necessary for its implementation via \Rightarrow on \mathcal{T} , the practicality of the interpretation \mathcal{HT} also depends on the closedness of \mathcal{T} under \Rightarrow . Namely, if the bisimulation collapse $G = \mathcal{HT}(\mathcal{G}) \Downarrow$ of the interpretation of some $\mathcal{G} \in \mathcal{H}$ were not contained in \mathcal{T} , then the converse interpretation \mathcal{TH} could not be applied to G in order to obtain the bisimulation collapse of \mathcal{G} .

A subclass \mathcal{K} of the term graphs over a signature Σ is called *closed under functional bisimulation* if, for all term graphs G, G' over Σ , whenever $G \in \mathcal{K}$ and $G \Rightarrow G'$, then also $G' \in \mathcal{K}$.

Note that for obtaining this property the use of variable backlinks, and backlinks for delimiter vertices is crucial (cf. Ex. 5.2).

Yet the class \mathcal{T} is actually not closed under \Rightarrow : See Fig. 9 at the top for a homomorphism from a non-eager-scope λ -term-graph to a term graph over $\Sigma_{S, \bullet}^\lambda$ that is not a λ -term-graph (as suggested by the overlapping scopes). The use of eager scope-closure remedies the situation, see at the bottom: then the bisimulation collapse is a λ -term-graph. This motivates the following theorem, which is proved in the extended report of [12]. It justifies property (P2) with \mathcal{T}_{eag} for \mathcal{T} .




Figure 9. \mathcal{T} is not closed under functional bisimulation, yet \mathcal{T}_{eag} is.

Theorem 5.8. The class \mathcal{T}_{eag} of eager-scope λ -term-graphs is closed under functional bisimulation \Rightarrow .

5.3 λ -term-graph semantics for λ_{letrec} -terms

We will consider in fact two interpretations of λ_{letrec} -terms as λ -term-graphs: first we define $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$ as the composition of $\llbracket \cdot \rrbracket_{\mathcal{H}}$ and \mathcal{HT} ; then we define the semantics $\llbracket \cdot \rrbracket_{\mathcal{T}}$ with more fine-grained S-sharing, which is necessary for defining a readback with the property (P3).

By composing the interpretation \mathcal{HT} of λ -ho-term-graphs as λ -term-graphs with the λ -ho-term-graph semantics $\llbracket \cdot \rrbracket_{\mathcal{H}}$, a semantics of λ_{letrec} -terms as λ -term-graphs is obtained. There is, however, a more direct way to define this semantics: by using an adaptation of the translation rules \mathcal{R} in Fig. 4, on which $\llbracket \cdot \rrbracket_{\mathcal{H}}$ is based. For this, let \mathcal{R}_S be the result of replacing the rule S in \mathcal{R} by the version in Fig. 10. While applications of this variant of the S-rule also shorten the abstraction-prefix, they additionally produce a delimiter vertex.




Figure 10. Delimiter-vertex producing version of the S-rule in Fig. 4

Here, at the end of the translation process, every loop on an indirection vertex with a prefix of length n is replaced by a chain of n S-vertices followed by a black hole vertex. Note that, while the system \mathcal{R}_S inherits all of the non-determinism of \mathcal{R} , the possible degrees of freedom have additional impact on the result, because now they also determine the precise degree of S-vertex sharing.

By analogous stipulations as in Def. 4.12 we define the conditions under which a λ -term-graph is called *eager-scope \mathcal{R}_S -generated*, or \mathcal{R}_S -generated *with minimal prefixes*, from a λ_{letrec} -term. For these notions, statements entirely analogous to Prop. 4.13 hold.

Definition 5.9. The semantics $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$ for λ_{letrec} -terms as λ -term-graphs is defined as $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min} : \text{Ter}(\lambda_{\text{letrec}}) \rightarrow \mathcal{T}_{\text{eag}}$, $L \mapsto \llbracket L \rrbracket_{\mathcal{T}}^{\min} :=$ the eager-scope term graph that is \mathcal{R}_S -generated with minimal prefixes from a garbage-free version L' of L .

For an example, see Ex. 5.14 below. In $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$, ‘min’ also indicates that λ -term-graphs obtained via this semantics exhibit minimal (in fact no) sharing (two or more incoming edges) of S-vertices. This is substantiated by the next proposition, in the light of the fact that \mathcal{HT} does not create any shared S-vertices.

Proposition 5.10. $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min} = \mathcal{HT} \circ \llbracket \cdot \rrbracket_{\mathcal{H}}$.

Hence $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$ only yields λ -term-graphs without sharing of S-vertices, and therefore its image cannot be all of \mathcal{T}_{eag} . As a consequence, we cannot hope to define a readback function rb with respect to $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$ that has the desired property (P3), because that requires that the image of the semantics is \mathcal{T}_{eag} in its entirety.

Therefore we modify the definition of $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$ to obtain another λ -term-graph semantics $\llbracket \cdot \rrbracket_{\mathcal{T}}$ with image $\text{im}(\llbracket \cdot \rrbracket_{\mathcal{T}}) = \mathcal{T}_{\text{eag}}$. This is achieved by letting the let-binding-structure of the λ_{letrec} -term influence the degree of S-sharing as much as possible, while staying eager-scope.

We say that a λ -ho-term-graph \mathcal{G} is *eager-scope \mathcal{R} -generated with maximal prefixes* from a λ_{letrec} -term L if \mathcal{G} is \mathcal{R} -generated from L by a translation process in which in applications of the let-rule the prefixes are chosen maximally, but so that the eager-scope property of the process is not compromised. It can be shown that this condition fixes the prefix lengths per application of the let-rule.

Definition 5.11. The semantics $\llbracket \cdot \rrbracket_{\mathcal{T}}$ for λ_{letrec} -terms as λ -term-graphs is defined as $\llbracket \cdot \rrbracket_{\mathcal{T}} : \text{Ter}(\lambda_{\text{letrec}}) \rightarrow \mathcal{T}_{\text{eag}}$, $L \mapsto \llbracket L \rrbracket_{\mathcal{T}} :=$




Figure 11. Translation of the λ_{letrec} -terms from Ex. 5.14 with the semantics $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$ and $\llbracket \cdot \rrbracket_{\mathcal{T}}$. For legibility some backlinks are merged.

the λ -term-graph that is eager-scope \mathcal{R}_S -generated with maximal prefixes from a garbage-free version L' of L .

Proposition 5.12. $\llbracket L \rrbracket_{\mathcal{T}}^{\min} \xrightarrow{S} \llbracket L \rrbracket_{\mathcal{T}}$ holds for all λ_{letrec} -terms L .

Now due to this, and due to Prop. 5.7, (iii), the statement of Thm. 4.15 can be transferred to \mathcal{T} , yielding property (P1) for $\llbracket \cdot \rrbracket_{\mathcal{T}}$.

Theorem 5.13. For all λ_{letrec} -terms L_1 and L_2 the following holds: $\llbracket L_1 \rrbracket_{\lambda^\infty} = \llbracket L_2 \rrbracket_{\lambda^\infty}$ if and only if $\llbracket L_1 \rrbracket_{\mathcal{T}} \leftrightarrow \llbracket L_2 \rrbracket_{\mathcal{T}}$.

Example 5.14. Consider the following four λ_{letrec} -terms:

$$\begin{aligned} L_1 &= \text{let } I = \lambda z. z \text{ in } \lambda x. \lambda y. \text{let } f = x \text{ in } ((y I) (I y)) (f f) \\ L_2 &= \lambda x. \text{let } I = \lambda z. z \text{ in } \lambda y. \text{let } f = x \text{ in } ((y I) (I y)) (f f) \\ L_3 &= \lambda x. \lambda y. \text{let } I = \lambda z. z, f = x \text{ in } ((y I) (I y)) (f f) \\ L'_3 &= \lambda x. \text{let } I = \lambda z. z \text{ in } \lambda y. \text{let } f = x, g = I \text{ in } ((y g) (g y)) (f f) \end{aligned}$$

The three possible fillings of the dashed area in Fig. 11 depict the translations $\llbracket L_1 \rrbracket_{\mathcal{T}}$, $\llbracket L_2 \rrbracket_{\mathcal{T}}$, and $\llbracket L_3 \rrbracket_{\mathcal{T}} = \llbracket L'_3 \rrbracket_{\mathcal{T}}$. The translations of the four terms with $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$ are identical:

$$\llbracket L_1 \rrbracket_{\mathcal{T}}^{\min} = \llbracket L_2 \rrbracket_{\mathcal{T}}^{\min} = \llbracket L_3 \rrbracket_{\mathcal{T}}^{\min} = \llbracket L'_3 \rrbracket_{\mathcal{T}}^{\min} = \llbracket L_1 \rrbracket_{\mathcal{T}}.$$

6. Readback of λ -term-graphs

In this section we describe how from a given λ -term-graph G a λ_{letrec} -term L that represents G (i.e. for which $\llbracket L \rrbracket_{\mathcal{T}} = G$ holds) can be ‘read back’. For this purpose we define a process based on term synthesis rules. It defines a readback function from λ -term-graphs to λ_{letrec} -terms. We illustrate this process by an example, formulate its most important properties, and sketch the proof of (P3).

The idea underlying the definition of the readback procedure is the following: For a given λ -term-graph G , a spanning tree T for G (augmented with a dedicated root node) is constructed that severs cycles of G at (some) recursive bindings or variable S-backlinks. Now the spanning tree T facilitates an inductive bottom-up (from the leaves upwards) synthesis process along T , which labels the edges of G (except for variable backlinks) with prefixed λ_{letrec} -terms. For this process we use local rules (see Fig. 13) that synthesise labels for incoming edges of a vertex from the labels of its outgoing edges. Eventually the readback of G is obtained as the label for the edge that singles out the root of term graph.

The design of the readback rules is based on a decision about where let-bindings are placed in the synthesised term. Namely there exists some freedom for these placements, as certain kind of shifts of let-expressions (let-floating steps [15]) preserve the λ -term-graph

Note that firing of the rules in step (Rb-5) of the readback process proceeds in bottom-up direction in the spanning tree, starting from the back-edges, with some room for parallelism concerning work in different subtrees. Furthermore observe that on all directed edges e (spanning tree edges or back edges) the rule applied to derive the edge label is uniquely determined by (is tied to) the label of the target vertex v of e , with the single exception of v being an indirection vertex. In that case one of the two indirection vertex rules applies, depending on whether e is a spanning-tree edge or a back-edge.


Proposition 6.2. Let G be a λ -term-graph. The process described in Def. 6.1 produces a complete edge labelling of the (modified) term graph, with label $(\ast[\cdot])L$ for the topmost edge, where L is a λ_{letrec} -term. Hence it yields L as the readback of G . Thus Def. 6.1 defines a function $\text{rb} : \mathcal{T} \rightarrow \text{Ter}(\lambda_{\text{letrec}})$, the *readback function*.

Example 6.3. See Fig. 12 for the illustration of the synthesis of the readback from an exemplary λ -term-graph. Full line edges are in the spanning tree, dotted line edges are not. Note that at the top vertex, no empty let-binding is created since the variant of the \top -vertex rule for empty binding groups is applied.

The following theorem validates property (P3), with \mathcal{T}_{eag} for \mathcal{T} .

Theorem 6.4. For all $G \in \mathcal{T}_{\text{eag}}$: $(\llbracket \cdot \rrbracket_{\mathcal{T}} \circ \text{rb})(G) = \llbracket \text{rb}(G) \rrbracket_{\mathcal{T}} \simeq G$, i.e., rb is a right-inverse of $\llbracket \cdot \rrbracket_{\mathcal{T}}$, and $\llbracket \cdot \rrbracket_{\mathcal{T}}$ a left-inverse of rb , up to \simeq . Hence rb is injective, and $\llbracket \cdot \rrbracket_{\mathcal{T}}$ is surjective, thus $\text{im}(\llbracket \cdot \rrbracket_{\mathcal{T}}) = \mathcal{T}_{\text{eag}}$.

Sketch of the Proof. Graph translation steps can be linked with corresponding readback steps in order to establish that the former roughly reverse the latter. Roughly, because e.g. reversing a λ -readback step necessitates both a λ - and a let-translation step. However, this holds only for a modification of the translation rules \mathcal{R}_5 from Fig. 4, Fig. 10 where the rules let (for let-expressions) and f (for occurrences of recursion variables) are replaced by the locally-operating versions in Fig. 14, and a initiating rule:



for creating a top vertex is added. Now the translation of a let-expression does no longer directly spawn translations of the bindings, but the bindings will only be translated later once their calls have been reached during the translation process of the in-part, or of the definitions of other already translated bindings. Note that in the let-rule in Fig. 14 function bindings are associated with the rightmost variable in the prefix, which corresponds to choosing $l_i = n$ in the let-rule in Fig. 4. While such a stipulation does not guarantee the eager-scope translation of every term, it actually does so for all λ_{letrec} -terms that are obtained by the readback

Please find in Fig. 16 on page 15 and in Fig. 17) on page 16 graphical arguments for the stepwise reversal of readback steps through translation steps. This establishes that graph translations steps reverse readback steps, which is the crucial step in the proof of the theorem. The proof uses induction on access paths, and an invariant that relates the eager-scope property localised for a vertex v with the applicability of the S-rule to the readback term synthesised at v . \square

7. Complexity analysis

Here we report on a complexity analysis for the individual operations from the previous sections, for the used standard algorithms, and overall, for compactification and unfolding-equivalence.

In the lemma below, (ii) and (v) justify the property (P4) of our methods. Items (iii) and (iv) detail the complexity of standard

methods when used for computing bisimulation collapse and bisimilarity of λ -term-graphs. For this note that first-order term graphs can be modelled by deterministic process graphs, and hence by DFAs. Therefore bisimilarity of term graphs can be computed via language equivalence of corresponding DFAs [20] (in time $O(n\alpha(n))$) [26], where α is the quasi-constant *inverse Ackermann function*) and bisimulation collapse through state minimisation of DFAs (in time $O(n \log n)$) [19].

- Lemma 7.1.** (i) $\text{size}(\llbracket L \rrbracket_{\mathcal{T}}) \in O(|L|^2)$ for $L \in \text{Ter}(\lambda_{\text{letrec}})$.
(ii) Translating $L \in \text{Ter}(\lambda_{\text{letrec}})$ into $\llbracket L \rrbracket_{\mathcal{T}} \in \mathcal{T}$ takes time $O(|L|^2)$.
(iii) Collapsing $G \in \mathcal{T}$ to $G \Downarrow$ is in $O(\text{size}(G) \log \text{size}(G))$.
(iv) Deciding bisimilarity of $G_1, G_2 \in \mathcal{T}$ requires time $O(n\alpha(n))$ for $n = \max\{\text{size}(G_1), \text{size}(G_2)\}$.
(v) Computing the readback $\text{rb}(G)$ for a given $G \in \mathcal{T}$ requires time $O(n \log n)$, for $n = \text{size}(G)$.

See Fig. 18 for an example that the size bound in item (i) of the lemma is tight.

Proposition 7.2. $\text{size}(\llbracket L \rrbracket_{\mathcal{T}}) \in O(|L|^2)$ for λ_{letrec} -terms L .

Based on this lemma, and on further considerations, we obtain the following complexity statements for our methods.

- Theorem 7.3.** (i) The computation for a λ_{letrec} -term L with $|L| = n$, of a maximally compactified form $(\text{rb} \circ \Downarrow \circ \llbracket \cdot \rrbracket_{\mathcal{T}})(L)$ of a λ_{letrec} -term L requires time $O(n^2 \log n)$. By using an S-unsharing operation unsh_S , a (typically smaller) λ_{letrec} -term $(\text{rb} \circ \text{unsh}_S \circ \Downarrow \circ \llbracket \cdot \rrbracket_{\mathcal{T}})(L)$ of size $O(n \log n)$ can be obtained, with the same time complexity.
(ii) The decision of whether two λ_{letrec} -terms L_1 and L_2 are unfolding equivalent requires time $O(n^2 \alpha(n))$ for $n = \max\{|L_1|, |L_2|\}$.

8. Implementation

We have implemented our methods in Haskell using the Utrecht University Attribute Grammar System. The implementation is available at <http://hackage.haskell.org/package/maxsharing/>. Output produced for three examples from this paper, and explanations for it, can be found in Appendix B; for all examples in [16].

9. Modifications, extensions and applications

We have described an adaptation of the bisimulation proof method for λ_{letrec} -terms. Recognising unfolding equivalence and increasing sharing are reduced to problems involving first-order term graphs. The principal idea is to use the nested scope structure of higher-order terms for an interpretation by term graphs with scope delimiters.

We conclude by describing easy modifications, rather direct extensions, and finally, promising areas of application for our methods.

9.1 Modifications

Implicit sharing of λ -variables. Multiple occurrences of the same λ -variable in a λ_{letrec} -term L are not shared (represented by a shared variable vertex) in the graph interpretation $\llbracket L \rrbracket_{\mathcal{H}}$. Consequently, our method compactifies the term $\lambda x. x x$ into $\lambda x. \text{let } f = x \text{ in } f f$. Such explicit sharing of variables is excessive for many applications. It can be remedied easily, namely by unsharing variable vertices before applying the readback, or by preventing the readback from introducing let-bindings when only a variable vertex is shared.

Avoiding aliases produced by the readback. The readback function in Section 6 is sensitive to the degree of sharing of S-vertices in the given λ -term-graph: it maps two λ -term-graphs that only differ in what concerns sharing of S-vertices to different λ_{letrec} -terms. Typically, for λ -term-graphs with maximal sharing of S-vertices this can produce let-bindings that are just ‘aliases’, such as g is alias




Figure 14. Modification of (two of) the translation rules in Fig. 4 for a variant definition of the λ -term-graph interpretation of λ_{letrec} -terms. Here the translation of a let-expression does not directly spawn translations for the binding equations, but the in-part has to be translated first.

$$\begin{array}{c}
 \frac{}{(\bar{p} x^v[\dots, f^w = L, \dots]) f} \text{ rec} \qquad \frac{(\bar{p} x^v[\dots, f^w = L, \dots]) L}{(\bar{p} x^v[\dots, f = L, \dots]) f} \text{ rec (name } w \text{ is fresh)} \\
 \\
 \frac{(\bar{p} x^v[B, f_1 = L_1, \dots, f_k = L_k]) L_0}{(\bar{p} x^v[B]) \text{ let } f_1 = L_1, \dots, f_k = L_k \text{ in } L_0} \text{ let}
 \end{array}$$

Figure 15. Formulation of the local translation rules in Fig. 14 in the form of inference rules.

for I in L'_3 from Ex. 5.14. This can be avoided in two ways: by slightly adapting the readback function, or by performing maximal unsharing of S-vertices before applying the readback as defined.

Preventing disadvantageous sharing. Introducing sharing at compile-time can cause ‘space leaks’, i.e. a needlessly high memory footprint, at run-time, because ‘a large data structure becomes shared [...], and therefore its space which before was reclaimed by garbage collection now cannot be reclaimed until its last reference is used’ [9]. For this reason, realisations of CSE [6] restrict the locally operating rewrite rules employed for introducing sharing by suitable conditions that account for the type of potentially shared subexpressions, and their strictness in the program. For our global method of introducing sharing via the bisimulation collapse, a different approach is needed.

Here the bisimulation collapse can be restricted so that sharing is not introduced at vertices that should not be shared. More precisely, it can be prevented that any unshared vertex (in-degree one) from a pre-determined set of ‘sharing-unfit’ vertices would have a shared vertex (in-degree greater than one) as its image in the bisimulation collapse. This can be achieved by modifying the graph interpretation $\llbracket \cdot \rrbracket_{\mathcal{T}}$. Any set of sharing-unfit positions in L gives rise to a set of sharing-unfit vertices in $\llbracket L \rrbracket_{\mathcal{T}}$. In the modification of $\llbracket L \rrbracket_{\mathcal{T}}$, special back-links are added from every sharing-unfit vertex with in-degree one to its immediate successor. These back-links prevent that such a sharing-unfit vertex v can collapse with another vertex v' without that also the predecessors of v and v' would collapse as well.

A more general notion of readback. Condition (P3) is rather rigorous in that it imposes sharing structure on λ_{letrec} that is specific to λ -term-graphs (degrees of S-sharing). For a weaker version of (P3) with \Leftarrow^S in place of isomorphism, a readback does not have to be injective, and, independently of how much S-sharing a translation into λ -term-graphs introduces, a readback function always exists.

9.2 Extensions

Full functional languages. In order to support programming languages that are based on λ_{letrec} like Haskell, additional language constructs need to be supported. Such languages can typically be desugared into a core language, which comprises only a small subset of language constructs such as constructors, case statements, and primitives. These constructs can be represented in an extension of λ_{letrec} by additional function symbols. In conjunction with a desugarer our methods are applicable to full programming languages.

Other programming languages, and calculi with binding constructs. Most programming languages feature constructs for grouping definitions that are similar to letrec. We therefore expect that our methods can be adapted to many imperative languages in particular, and may turn out to be fruitful for optimising compilers. Our methods for achieving maximal sharing certainly generalise to theoretical frameworks, and calculi with binding constructs, such as the π -calculus [24], and higher-order rewrite systems (e.g. CRSs and HRSs, [32]) as used here for the formalisation of λ_{letrec} .

Fully-lazy lambda-lifting. There is a close connection between our methods and fully-lazy lambda-lifting [21, 29]. In particular, the required-variable and scope analysis of a λ_{letrec} -term L on which the λ -term-graph-translation $\llbracket L \rrbracket_{\mathcal{T}}$ is based is closely analogous to the one needed for extracting from L the supercombinators in the result \hat{L} of fully-lazy lambda-lifting L . Moreover, the fully-lazy lambda-lifting transformation can even be implemented in a natural way on the basis of our methods. Namely as the composition $\text{rb}_{LL} \circ \llbracket \cdot \rrbracket_{\mathcal{T}}$ of the translation $\llbracket \cdot \rrbracket_{\mathcal{T}}$ into λ -term-graphs, where rb_{LL} is a variant readback function that, for a given λ -term-graph, synthesises the system \hat{L} of supercombinators, instead of the λ_{letrec} -term $\text{rb}(L)$.

Maximal sharing on supercombinator translations of λ_{letrec} -terms. λ_{letrec} -terms L correspond to supercombinator systems \hat{L} , the result of fully-lazy lambda-lifting L : the combinators in \hat{L} correspond to ‘extended scopes’ [11] (or ‘skeletons’ [2]) in L , and supercombinator reduction steps on \hat{L} correspond to weak β -reduction steps L . In the case of λ -calculus this has been established by Balabonski [2]. Via this correspondence the maximal-sharing method for λ_{letrec} -terms can be lifted to obtain a maximal-sharing method systems of supercombinators obtained by fully-lazy lambda-lifting.

Non-eager scope-closure strategies. We focused on eager-scope translations, because they facilitate maximal sharing, and guarantee that interpretations of unfolding-equivalent λ_{letrec} -terms are bisimilar. Yet every scope-closure strategy [11] induces a translation and its own notion of maximal sharing. For adapting our maximal sharing method it is necessary to modify the translation into first-order term graphs in such a way that the image class obtained is closed under homomorphism (\mathcal{T} is not closed under \Rightarrow , unlike its subclass \mathcal{T}_{eag}). This can be achieved by using delimiter vertices also below variable vertices to close scopes that are still open [12, report].

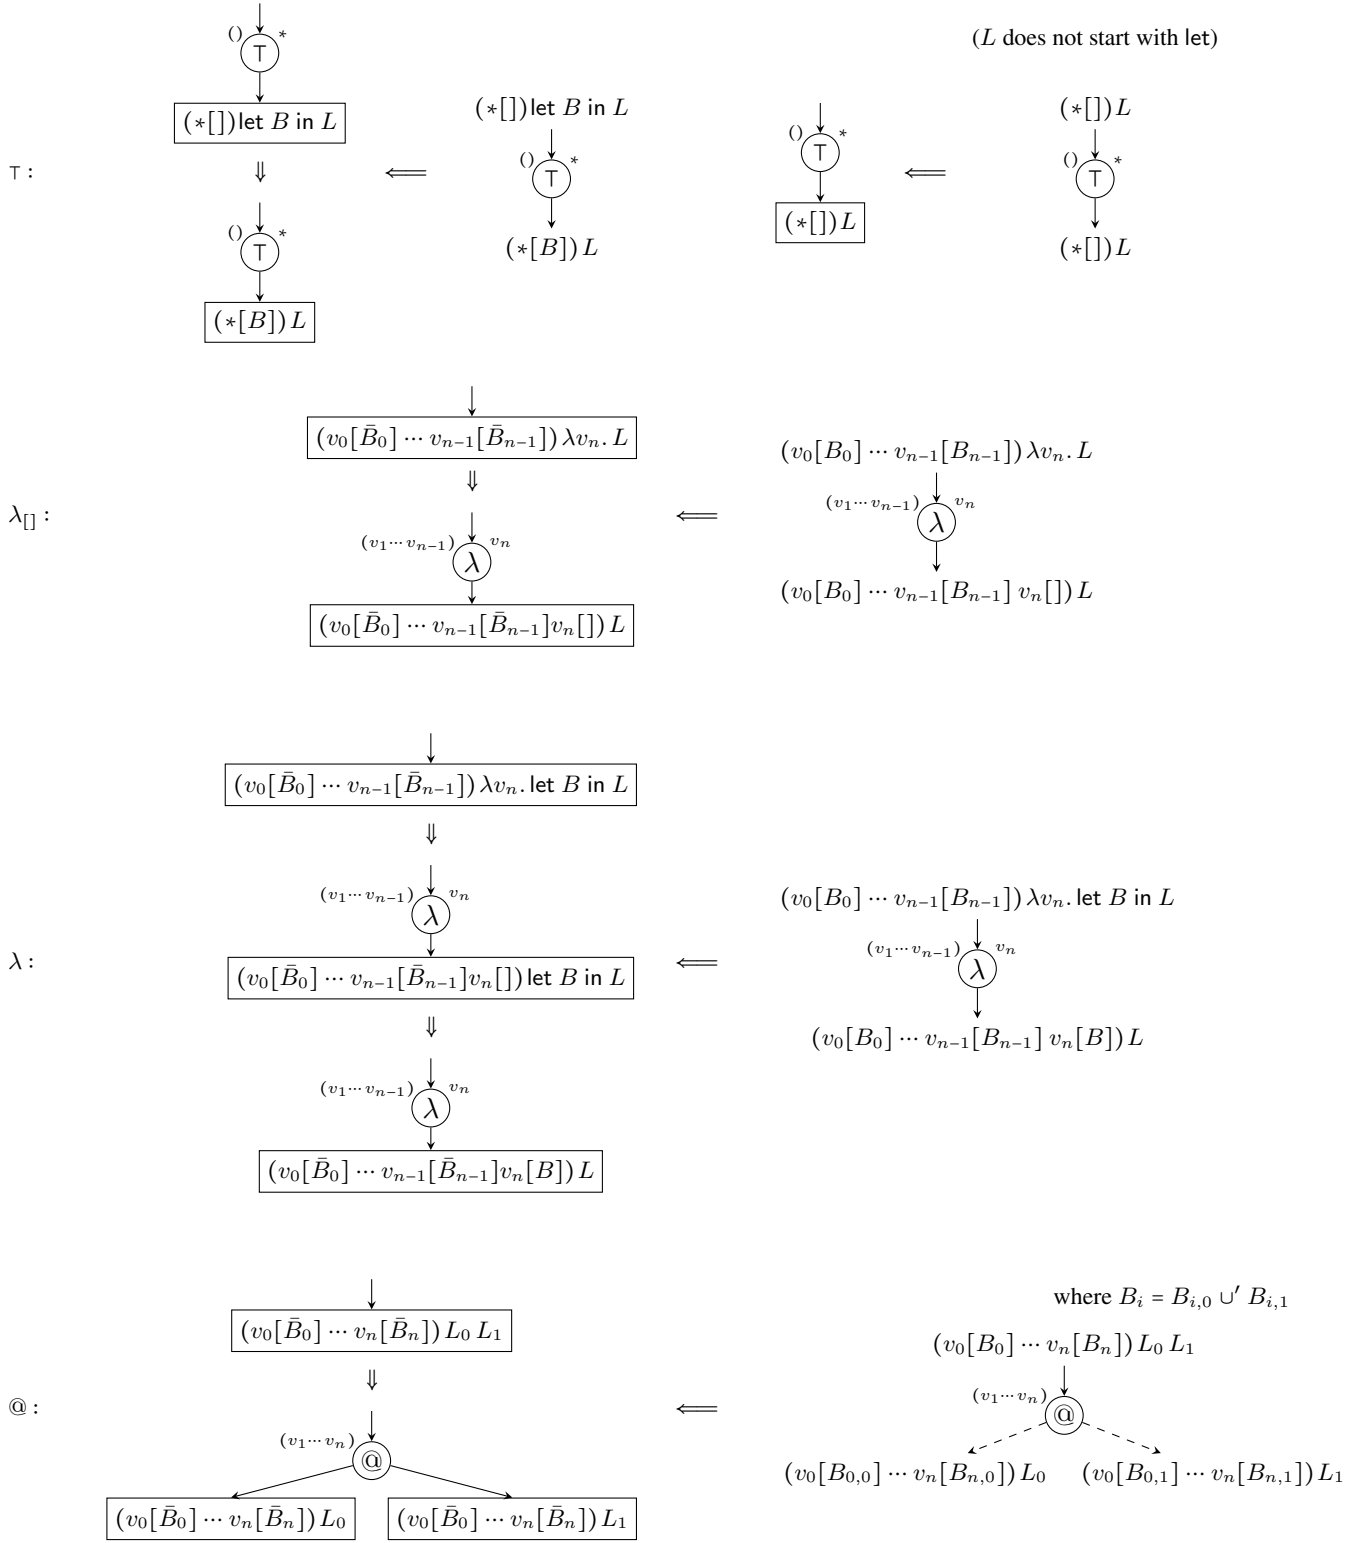


Figure 16. Reversal of the readback steps for top vertices, abstraction vertices, and application vertices by corresponding translation steps.



16

Consider the finite λ -terms M_n with n occurrences of bindings λx_2 :

$$\lambda x_0 x_1. x_0 x_1 \lambda x_2. x_0 x_1 \lambda x_1. x_0 x_2 \lambda x_2. x_0 x_1 \dots \lambda x_2. x_0 x_1 x_2$$

Then $|M_n| \in O(n)$. But both the transformation of M_n into a de-Brujin index representation:

$$\lambda(\lambda(@(@(\lambda(S(0), 0), \lambda(@(@(\lambda(S^2(0), S(0)), \lambda(@(@(\lambda(S^3(0), S(0)), \lambda(@(@(\lambda(S^4(0), S(0)), \dots \lambda(@(@(\lambda(S^{2n}(0), S(0)), 0))))))))))))))$$

and the rendering of M_n with respect to the eager scope-delimiting strategy:

$$\lambda(\lambda(@(@(\lambda(S(0), 0), \lambda(@(\lambda(S(@(\lambda^1(0), 0)), \lambda(@(\lambda(S(@(\lambda^2(0), 0)), \lambda(@(\lambda(S(@(\lambda^3(0), 0)), \dots \lambda(@(\lambda(S(@(\lambda^{2n-1}(0), 0)), 0))))))))))))))$$

have size $O(n^2)$.

Figure 18. Example of a sequence $\{M_n\}_n$ of finite λ -terms M_n whose translation into λ -term-graphs grows quadratically in the size of M_n .

Weaker notions of sharing. The presented methods deal with sharing as expressed by letrec that is horizontal, vertical, or twisted [4]. By contrast, the construct μ [4, 13] expresses only vertical, and the non-recursive let only horizontal, sharing. By restricting bisimulation, our methods can be adapted to the λ -calculus with μ , or with let.

Nested term graphs. The nested scope structure of a λ_{letrec} -term can also be represented by a nested structure of term graphs. The representation of a λ_{letrec} -term as a ‘nested term graph’ [17] starts with an ordinary term graph in which some of the vertices are labelled by ‘nested’ symbols that designate outermost bindings together with their scope. Any such vertex is additionally associated with a usual term graph that specifies the subterm context describing the scope, where any inner scopes are again expressed by nested symbols. The association between nested symbols and their term graph specifications is required to be tree-like. The implementation result developed here can be generalised to show that nested term graphs can be implemented faithfully by first-order term graphs [17].

9.3 Applications

Maximal sharing at run-time. Maximal sharing can be applied repeatedly at run-time in order to regain a maximally shared form, thereby speeding up evaluation. This is reminiscent of ‘collapsed tree rewriting’ [30] for evaluating first-order term graphs represented as maximally shared dags. Since the state of a program in the memory at run-time is typically represented as a supercombinator graph, compactification by bisimulation collapse can take place directly on that graph (see Sec. 9.2), no translation is needed. Compactification can be coupled with garbage collection as bisimulation collapse subsumes some of the work required for a mark and sweep garbage collector. However, a compromise needs to be found between the costs for the optimisation and the gained efficiency.

Compile-time optimisation phase. Increasing sharing facilitates potential gains in efficiency. Our method generalises common subexpression elimination, but therefore it also inherits its shortcomings: the cost of sharing (e.g. of very small functions) might exceed the gain. In non-strict functional languages, sharing can cause ‘memory leaks’ [6]. Therefore, similar as for CSE, additional dynamic analyses like binding-time analysis [28], and heuristics to restrict sharing in cases when it is disadvantageous [10, 29] are advisable.

Additional prevention of disadvantageous sharing. While static analysis methods for preventing sharing that may be disadvantageous at run-time can be adapted from CSE to the maximal-sharing method (see Sec. 9.1), this has yet to be investigated for binding-time analysis [28] and a sharing analysis of partial applications [10].

Code improvement. In programming it is generally desirable to avoid duplication of code. As extension of CSE, our method is able to detect code duplication. The bisimulation collapse of the term graph interpretation of a program can, together with the readback, provide guidance on how code can be refactored into a more compact form. This application requires some fine-tuning to avoid

excessive behaviour like the explicit sharing of variable occurrences (see Sec. 9.1). Yet for this only lightweight additional machinery is needed, such as size constraints or annotations to restrict the bisimulation collapse.

Function equivalence. Recognising whether two programs implement the same function is undecidable. Still, this problem is tackled by proof assistants, and by automated theorem provers used in type-checkers of compilers for dependently-typed programming languages such as Agda. For such systems co-inductive proofs are more difficult to find than inductive ones, and require more effort by the user. Our method for deciding unfolding-equivalence could help to develop new approaches to finding co-inductive proofs.

Acknowledgment We want to thank Vincent van Oostrom for extensive feedback on a draft, Doaitse Swierstra and Dimitri Hendriks for helpful comments, and Jeroen Keiren for a suggestion concerning restricting the bisimulation collapse. We also thank the anonymous reviewers for their comments, and a number of stimulating questions.

References

- [1] A. Asperti and S. Guerrini. *The Optimal Implementation of Functional Programming Languages*. Cambridge University Press, 1998.
- [2] T. Balabonski. A unified approach to fully lazy sharing. In *Proceedings of POPL '12*, pages 469–480, New York, NY, USA, 2012. ACM.
- [3] R. S. Bird and R. Patterson. de Bruijn notation as a nested datatype. *Journal of Functional Programming*, 9(1):77–91, 1999.
- [4] S. Blom. *Term Graph Rewriting – Syntax and Semantics*. PhD thesis, Vrije Universiteit Amsterdam, 2001.
- [5] M. v. d. Brand and P. Klint. ATERMs for manipulation and exchange of structured data: It’s all about sharing. *Information and Software Technology*, 49(1):55–64, 2007.
- [6] O. Chitil. Common Subexpressions Are Uncommon in Lazy Functional Languages. In *Selected Papers from the 9th International Workshop IFL (IFL '97)*, pages 53–71, London, UK, UK, 1998. Springer-Verlag.
- [7] O. Danvy and U. P. Schultz. Lambda-lifting in quadratic time. *Journal of Functional and Logic Programming*, 2004, 2004.
- [8] N. G. de Bruijn. Lambda Calculus Notation with Nameless Dummies, a Tool for Automatic Formula Manipulation, with Applic. to the Church-Rosser Theorem. *Indagationes Mathematicae*, 34:381–392, 1972.
- [9] A. L. de Medeiros Santos. *Compilation by Transformation in Non-Strict Functional Languages*. PhD thesis, University of Glasgow, 1995.
- [10] B. Goldberg and P. Hudak. Detecting Sharing of Partial Applications in Functional Programs. Technical Report YALEU/DCS/RR-526, Department of Computer Science, Yale University, March 1987.
- [11] C. Grabmayer and J. Rochel. Expressibility in the Lambda Calculus with Letrec. Technical report, arXiv, August 2012. arXiv:1208.2383.
- [12] C. Grabmayer and J. Rochel. Term Graph Representations for Cyclic Lambda Terms. In *Proceedings of TERMGRAPH 2013*, number 110 in EPTCS, 2013. For an extended report see: arXiv:1308.1034.


- [13] C. Grabmayer and J. Rochel. Expressibility in the Lambda Calculus with μ . In *Proceedings of RTA 2013*, 2013. Report: arXiv:1304.6284.
- [14] C. Grabmayer and J. Rochel. Expressibility in the Lambda Calculus with μ . Technical report, <http://arxiv.org>, 2013. Report extending [13], arXiv:1304.6284.
- [15] C. Grabmayer and J. Rochel. Confluent Let-Floating. In *Proceedings of IWC 2013 (2nd International Workshop on Confluence)*, 2013.
- [16] C. Grabmayer and J. Rochel. Maximal Sharing in the Lambda Calculus with letrec. Technical report, arXiv, January 2014. arXiv:1401.1460.
- [17] C. Grabmayer and V. van Oostrom. Nested Term Graphs. Technical report, arXiv, May 2014. arXiv:1405.6380.
- [18] D. Hendriks and V. van Oostrom. λ . In F. Baader, editor, *Proceedings CADE-19*, volume 2741 of *LNAI*, pages 136–150. Springer, 2003.
- [19] J. Hopcroft. An $n \log n$ Algorithm for Minimizing States in a Finite Automata. Technical report, Stanford University, CA, USA, 1971.
- [20] J. Hopcroft and R. Karp. A Linear Algorithm for Testing Equivalence of Finite Automata. Technical report, Cornell University, 1971.
- [21] R. Hughes. Supercombinators: A new implementation method for applicative languages. In *LFP '82: Proceedings of the 1982 ACM symposium on LISP and functional programming*, pages 1–10, 1982.
- [22] T. Johnsson. Lambda lifting: Transforming programs to recursive equations. In *FPCA*, pages 190–203, 1985.
- [23] J. Ketema and J. G. Simonsen. Infinitary Combinatory Reduction Systems. *Information and Computation*, 209(6):893 – 926, 2011.
- [24] R. Milner. *Communicating and mobile systems: the π -calculus*. Cambridge University Press, 1999.
- [25] M. T. Morazán and U. P. Schultz. Optimal lambda lifting in quadratic time. In *Workshop IAFL 2007*, number 5083 in LNCS. Springer, 2008.
- [26] D. A. Norton. Algorithms for Testing Equivalence of Finite Automata. Master's thesis, Dept. of Computer Science, Rochester Institute of Technology, 2009. <https://ritdml.rit.edu/handle/1850/8712>.
- [27] V. v. Oostrom, K.-J. van de Looij, and M. Zwitterlood. Lambdascope. Extended Abstract, Workshop ALPS, Kyoto, April 10th 2004, 2004.
- [28] J. Palsberg and M. Schwartzbach. Binding-time analysis: abstract interpretation versus type inference. In *Int. Conf. on Computer Languages*, 1994, pages 289–298, 1994.
- [29] S. L. Peyton Jones. *The Implementation of Functional Programming Languages*. Prentice-Hall, Inc., 1987.
- [30] D. Plump. *Evaluation of Functional Expressions by Hypergraph Rewriting*. PhD thesis, Universität Bremen, 1993.
- [31] J. Rochel and C. Grabmayer. Avoiding Repetitive Evaluation Patterns in the Lambda Calculus with Letrec (Work in Progress). In R. Echahed, editor, *Proceedings of TERMGRAPH 2011*, volume 48 of *EPTCS*, on arXiv:1102.2268, pages 85–100. arxiv.org, 2011.
- [32] Terese. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2003.
- [33] C. P. Wadsworth. *Semantics and Pragmatics of the Lambda-Calculus*. PhD thesis, University of Oxford, 1971.

A. Example for the translation λ_{letrec} -terms into λ -ho-term-graphs and into λ -term-graphs

For two terms from the paper we provide the stepwise translation of λ_{letrec} -terms into λ -ho-term-graphs.

We start off by a simple example, namely the translation of the term P from Example 1.2 on page 1. There is no application of the S-rule, thus it yields the same sequence of graphs regardless of

whether the rules \mathcal{R} from Fig. 4 are used, by which a λ -ho-term-graph is produced, or the modified rules \mathcal{R}_5 (the result of dropping the S-rule from \mathcal{R} , and using the S-rule in Fig. 10 instead), by which λ -term-graphs with an abstraction-prefix function are produced.

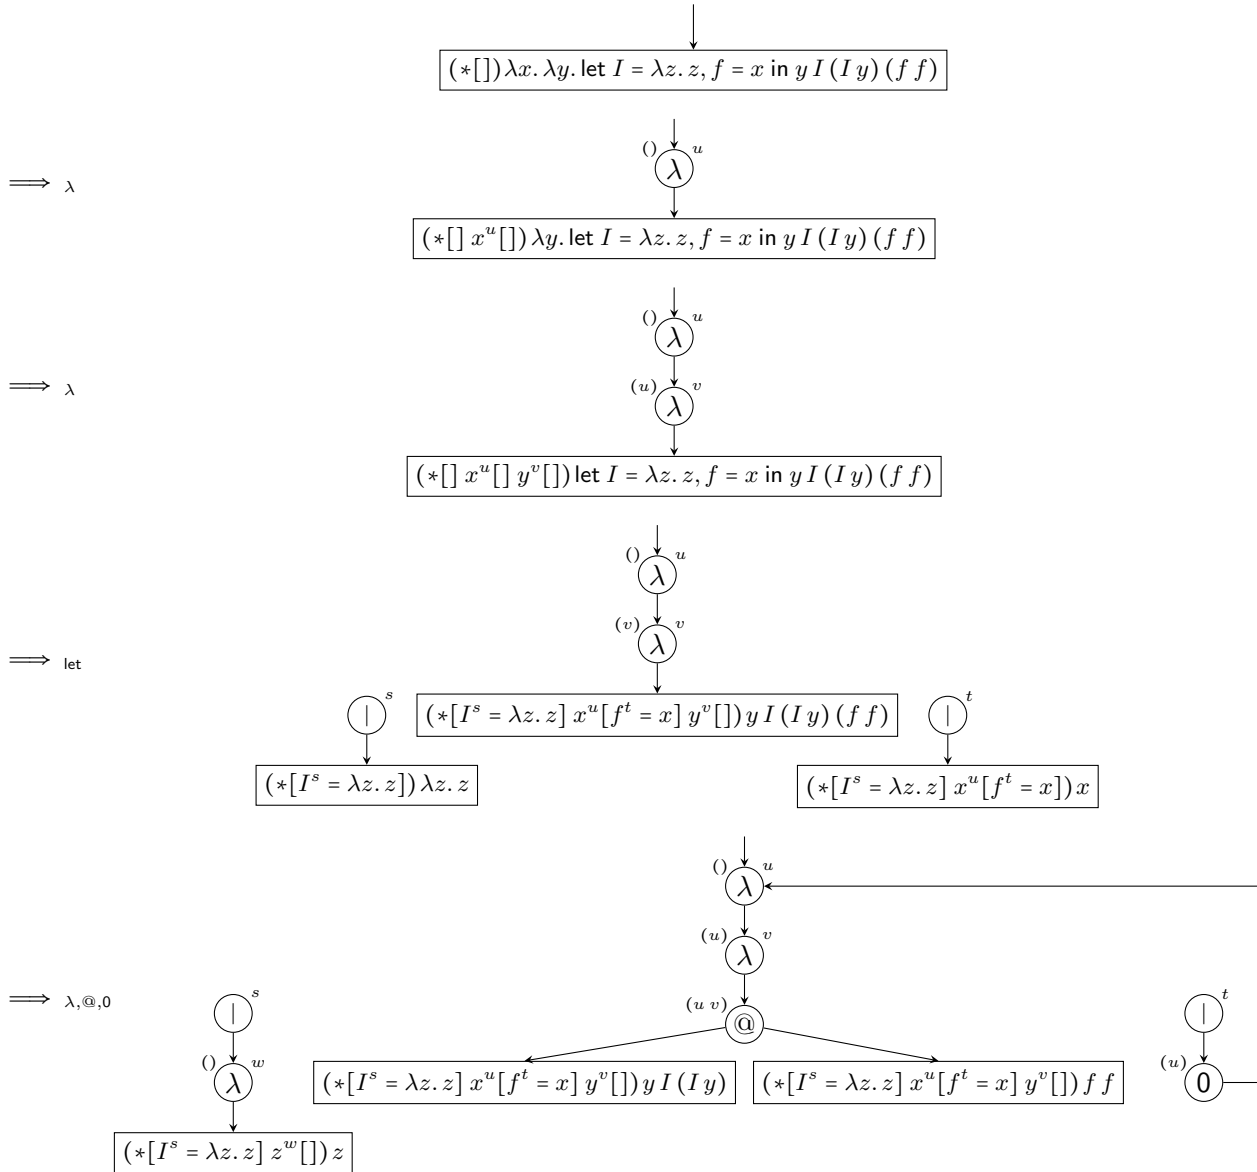



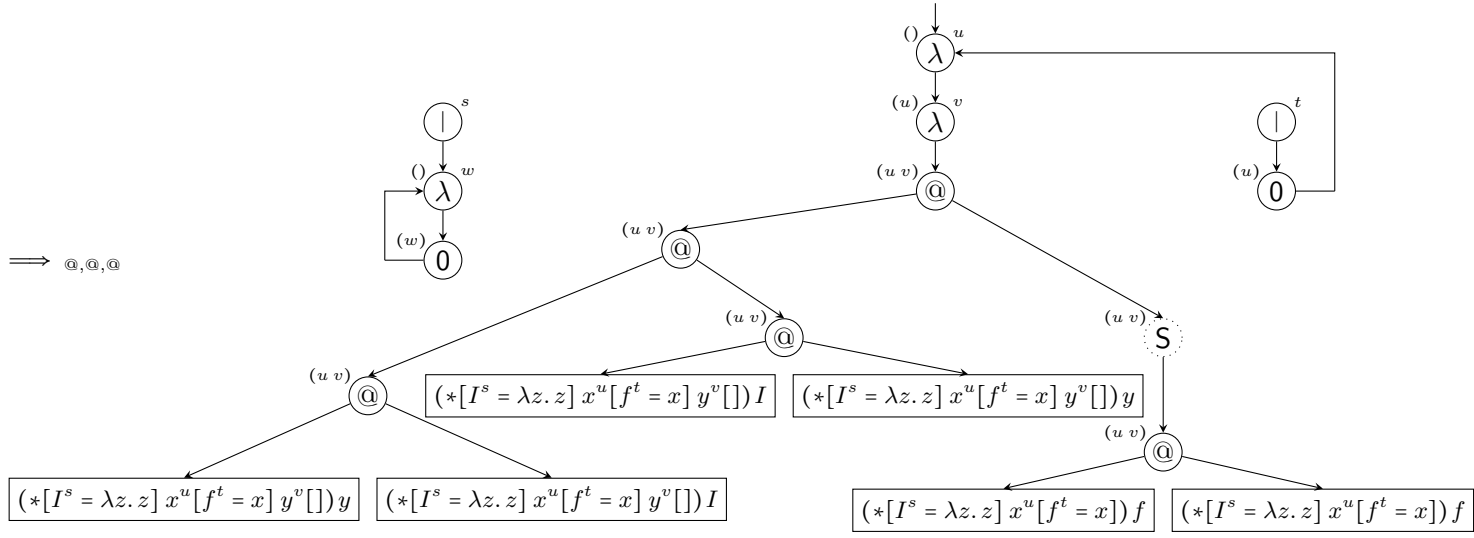
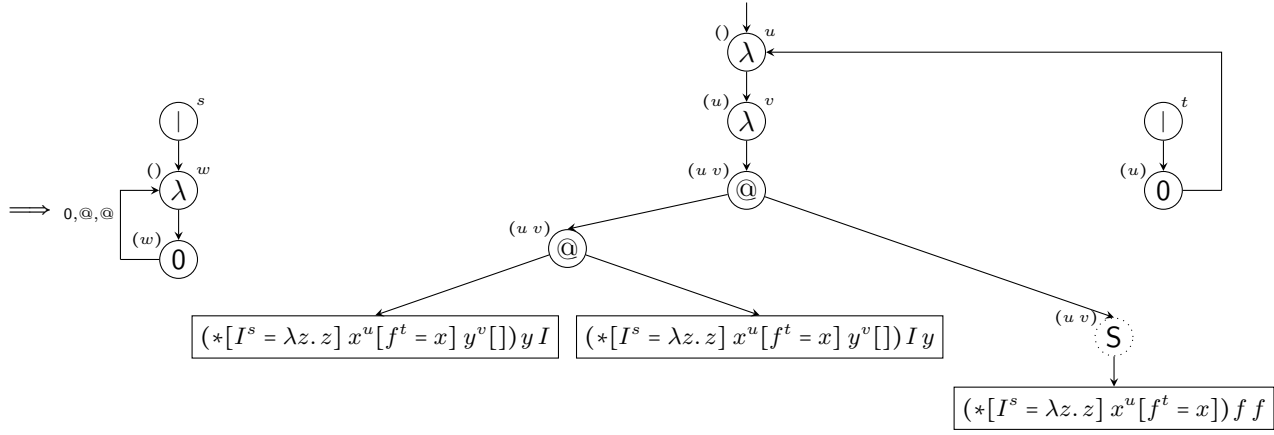
We continue with the term L_2 from Example 5.14 on page 11. We translate it in two different ways corresponding to the first-order term graph semantics $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$ from Def. 5.9 and $\llbracket \cdot \rrbracket_{\mathcal{T}}$ from Def. 5.11, respectively. Both sequences of translation steps yield the same λ -ho-term-graph $\llbracket L_2 \rrbracket_{\mathcal{H}}$, but obtain the different λ -term-graphs $\llbracket L_2 \rrbracket_{\mathcal{T}}^{\min}$ and $\llbracket L_2 \rrbracket_{\mathcal{T}}$ in Fig. 11.

These two stepwise translation processes exemplify, on the one hand (when ignoring the dotted S-vertices that do not occur in λ -ho-term-graphs) the translation into λ -ho-term-graphs according to the rules \mathcal{R} from Fig. 4, and on the other hand (now taking the dotted delimiter S-vertices into account) the translation into λ -term-graphs according to the modified rules \mathcal{R}_S .

In each step, one or more translation rules (whose names are indicated as subscripts in the steps) are applied to the translation boxes in the graph. When no more rules are applicable, indirection vertices are erased. Both translations are eager-scope (i.e. applications of S-rules are given priority, and prefixes lengths are chosen too be small enough in the let-rule, see Sec. 4.2) but differ in how they resolve the non-determinism due to different choices for the prefix lengths l_1, \dots, l_k in the let-rule.

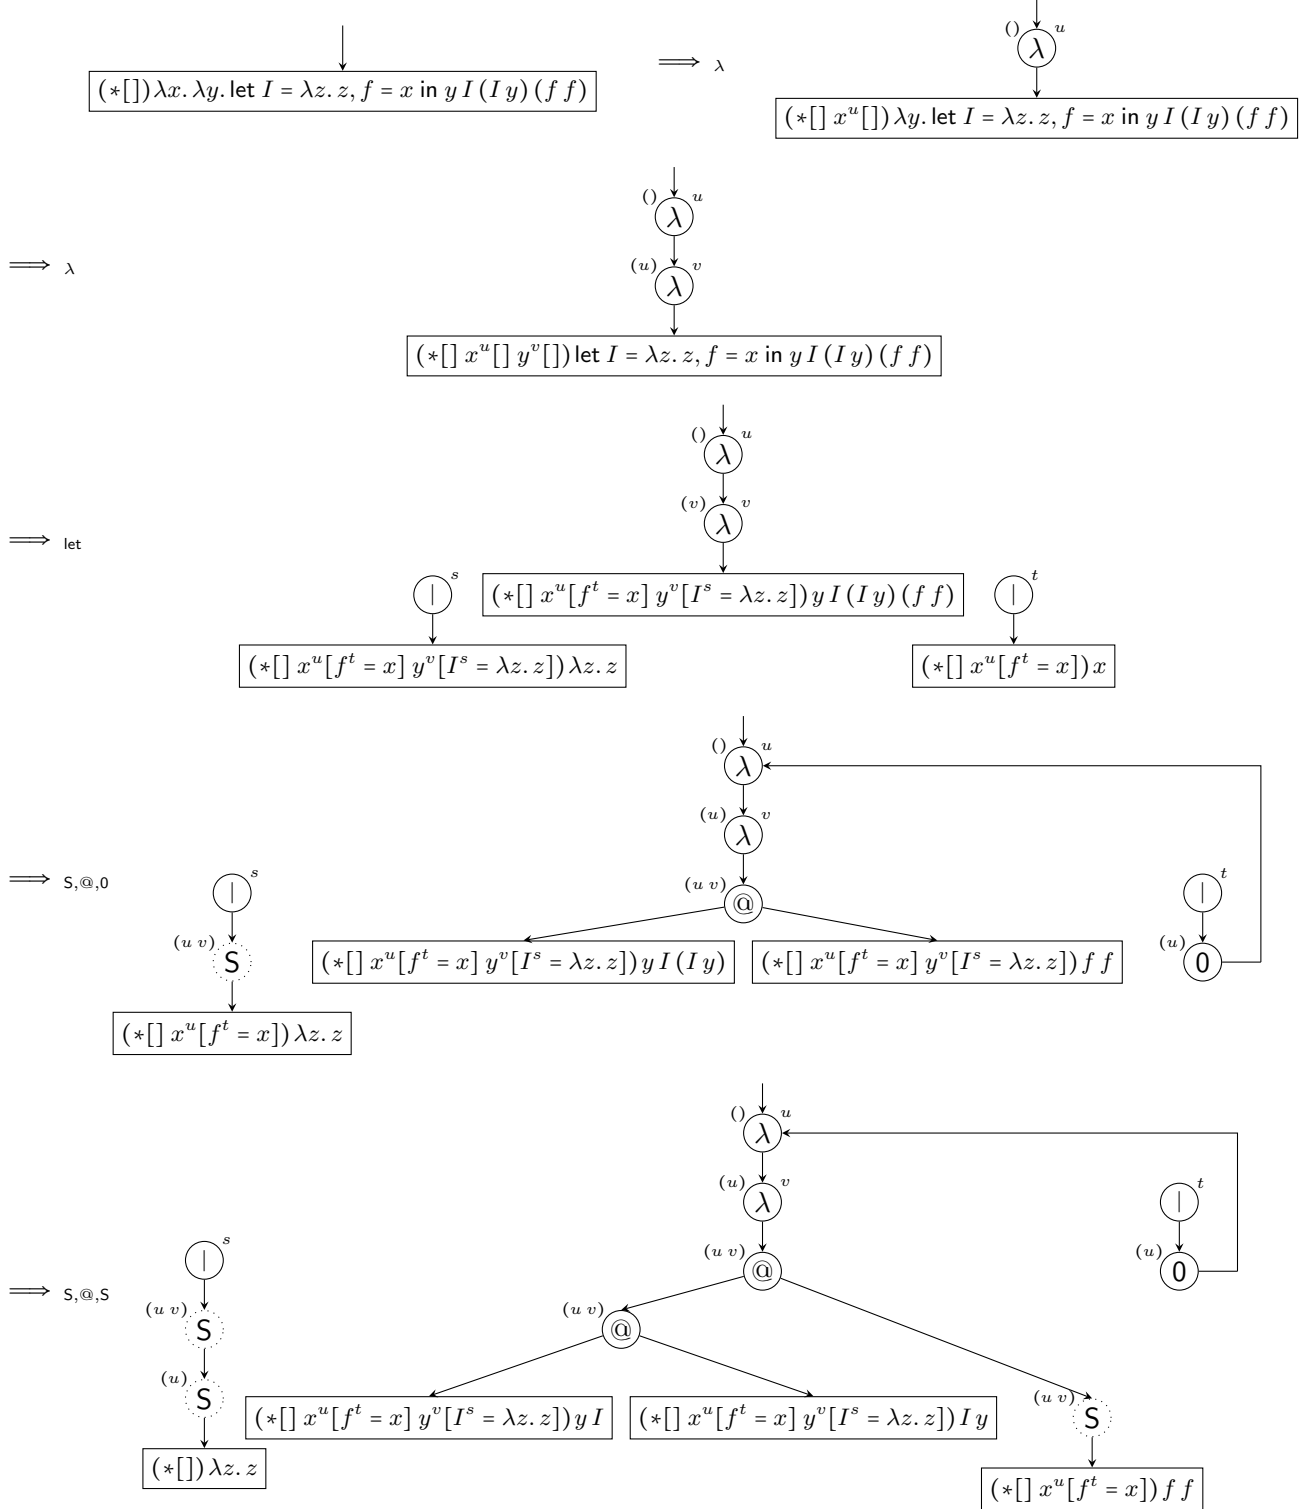
First we consider the translations of L_2 with $\llbracket \cdot \rrbracket_{\mathcal{H}}$ and $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$, i.e. the translation process in which prefix lengths are chosen minimally when applying the let-rule.



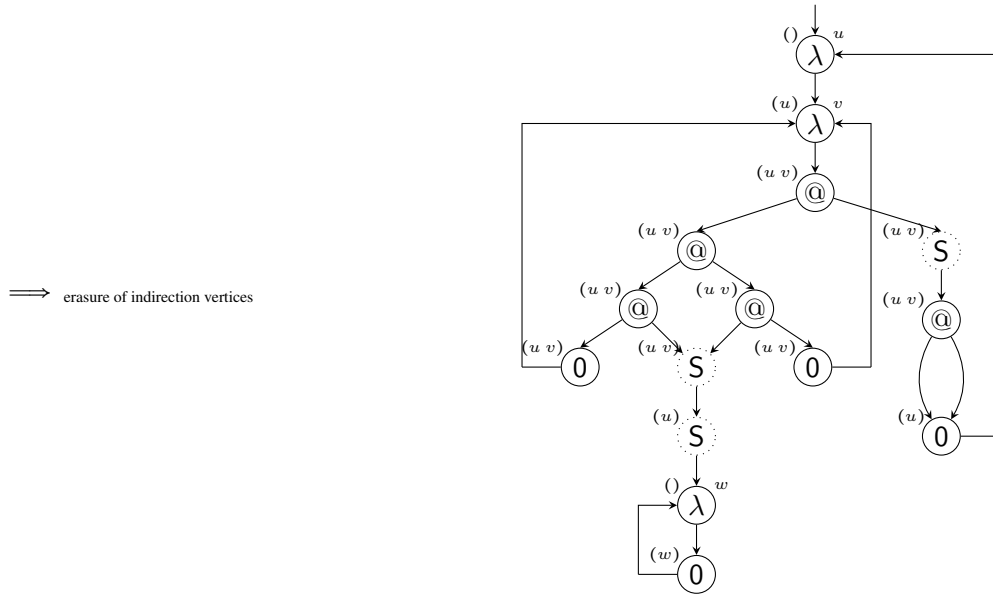

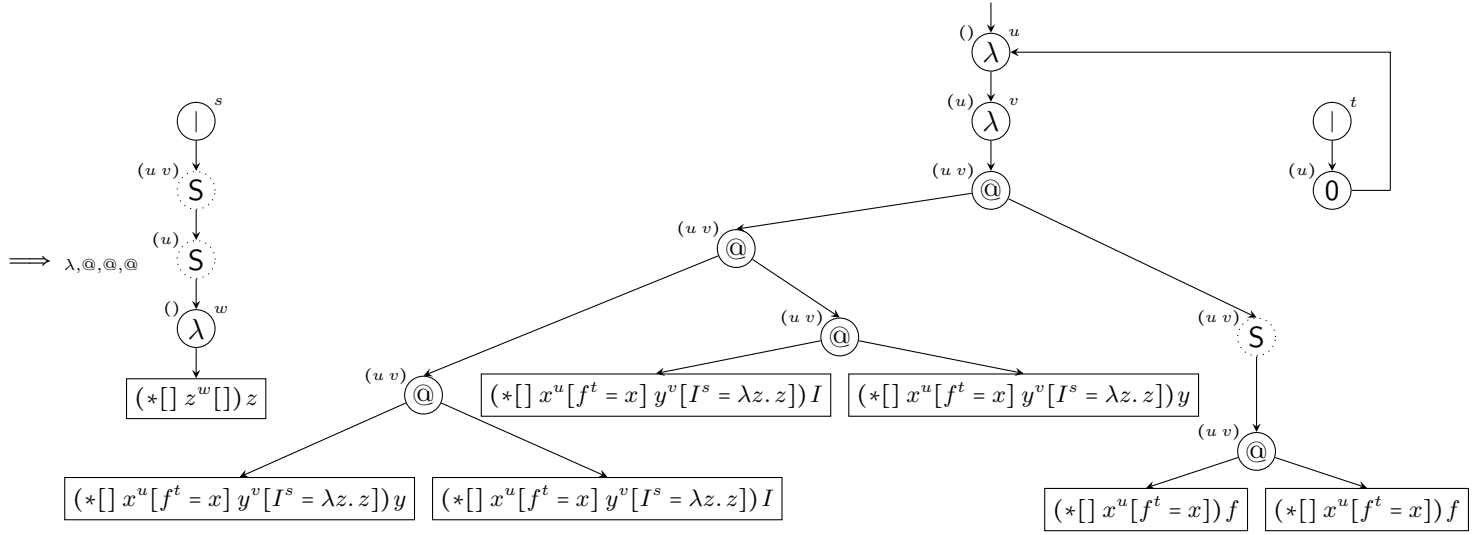


Second, we give the translation of the same term L_2 (from Example 5.14 on page 11) with the process needed for the first-order term graph semantics $\llbracket \cdot \rrbracket_{\mathcal{T}}$, yielding $\llbracket L_2 \rrbracket_{\mathcal{T}}$ in Fig. 4, and, at first sight incidentally⁴, the corresponding λ -ho-term-graph $\llbracket L_2 \rrbracket_{\mathcal{H}}$. Note that the resulting λ -ho-term-graph (ignore the dotted delimiter S-

vertices) is again $\llbracket L_2 \rrbracket_{\mathcal{H}}$, that is, it is identical⁴ with the one that was produced by the translation process above. Yet the obtained λ -term-graph (now taking the dotted S-vertices into account) $\llbracket L_2 \rrbracket_{\mathcal{T}}$ differs from the λ -term-graph $\llbracket L_2 \rrbracket_{\mathcal{T}}^{\min}$ obtained above by exhibiting a higher degree of S-sharing.



⁴ The fact that this is actually not just a coincidence in this specific example is an easy consequence of Prop. 5.7, (vi), Prop. 5.10, and Prop. 5.12.



B. Implementation Showcase

To demonstrate the realisability of our method, and for further illustration, we include the output of our implementation for the examples used in the paper. The implementation is called `maxsharing` and is available on Hackage. It is written in Haskell and therefore requires the Haskell Platform to be installed.

B.1 Example 1.1

λ -letrec-term:
 $(\lambda x. x) (\lambda x. x)$

translation used: minimal prefix lengths

scoped (with adbmals):
 $(\lambda x. x) (\lambda x. x)$

scoped (with scope delimiters and nameless abstractions):
 $(\lambda. 0) (\lambda. 0)$

derivation:

```

----- 0      ----- 0
(* x) x      (* x) x
-----  λ      -----  λ
(*) λx. x    (*) λx. x
-----  @
(*) (λx. x) (λx. x)


```

DFA: writing to file

minimised DFA: writing to file

spanning tree:
 $A(|F(L(0)), F)$

readback:
 $\text{let } F = \lambda x. x$
 $\text{in } F F$



the original term as recognised by the parser

The user can specify which term graph semantics shall be used in processing the term. All further output is with respect to that translation. The two options are:

- $\llbracket \cdot \rrbracket_{\mathcal{T}}$ is indicated by: maximal prefix lengths while maintaining eager scope-closure
- $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$ is indicated by: minimal prefix lengths.

For this example both semantics yield the same output.

term with scope delimiters; see Example 4.2 for a more interesting case and more explanations

Derivation according to the proof system in Fig. 6 that shows the translation as a stepwise process and includes all subterms with their abstraction prefixes. Note that in the notation of the prefixed terms the binding annotation of a variable is omitted if it is empty. Also, it only includes the names of the function variables, not their entire definition.

Even though the correspondence between the derivations in the proof system in Fig. 6 and the translations $\llbracket \cdot \rrbracket_{\mathcal{T}}$ and $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$ is not provided here, we think that the derivation can help as an illustration of the translation process and its result.

The implementation produces a graphical depiction for the term's λ -term-graph in DFA form, as well as for the minimised form of the DFA; see below for pictures.

A textual representation of the minimised DFA's spanning tree, used for readback. It is displayed in first-order term rewriting notation, i.e. with a unary function symbols L , S for abstraction and scope delimiters, a binary function symbol A for application and a nullary symbol 0 for abstraction variable occurrences. Furthermore there is a class of function symbols written as a vertical bar followed by an upper-case variable name $|F$, $|G$, etc. which signify a vertex with multiple incoming non-backlink edges, and therefore a vertex that will be the root of a shared subterm. There is also a class of corresponding nullary function symbols F , G , etc. which represent non-backlink, non-spanning-tree edges to these shared vertices.

Readback of the minimised DFA.

B.2 Example 1.2

Also for the terms L and P of Ex. 1.2 on page 1 the translations are identical for $\llbracket \cdot \rrbracket_{\mathcal{H}}$ and $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$.

λ -letrec-term:
 $\lambda f. \text{let } r = f \text{ } r$
in r

translation used: minimal prefix lengths

scoped (with adbmals):
 $\lambda f. \text{let } r = f \text{ } r$
in r

scoped (with scope delimiters and nameless abstractions):
 $\lambda. \text{let } r = 0 \text{ } r$
in r



derivation:
----- 0
(* f[r]) f (* f[r]) r
----- @
(* f[r]) f r (* f[r]) r
----- let
(* f) let r = f r
in r
----- λ
(*) $\lambda f. \text{let } r = f \text{ } r$
in r

DFA: writing to file

minimised DFA: writing to file

spanning tree:
 $L(|F(A(0, F)))$

readback:
 $\lambda x. \text{let } F = x \text{ } F$
in F



λ -letrec-term:
 $\lambda f. \text{let } r = f \text{ } (f \text{ } r)$
in r

translation used: minimal prefix lengths

scoped (with adbmals):
 $\lambda f. \text{let } r = f \text{ } (f \text{ } r)$
in r

scoped (with scope delimiters and nameless abstractions):
 $\lambda. \text{let } r = 0 \text{ } (0 \text{ } r)$
in r



derivation:
----- 0
(* f[r]) f (* f[r]) r
----- @
(* f[r]) f (* f[r]) f r
----- @
(* f[r]) f (f r) (* f[r]) r
----- let
(* f) let r = f (f r)
in r
----- λ
(*) $\lambda f. \text{let } r = f \text{ } (f \text{ } r)$
in r

DFA: writing to file

minimised DFA: writing to file

spanning tree:
 $L(|F(A(0, F)))$

readback:
 $\lambda x. \text{let } F = x \text{ } F$
in F



B.3 Example 4.2

Again, for the terms L_1 and L_2 of Ex. 4.2 on page 5 the translations are identical for $\llbracket \cdot \rrbracket_{\mathcal{H}}$ and $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$.

```
λ-letrec-term:
let f = λx. (λy. f y) x
in f
```

translation used: minimal prefix lengths

```
scoped (with adbmals):
let f = λx. /x. λy. /y. f y x
in f
```

```
scoped (with scope delimiters and nameless abstractions):
let f = λ. S((λ. S(f) 0)) 0
in f
```



```
derivation:
(*[f]) f
----- S ----- 0
(*[f] y) f    (*[f] y) y
----- @
(*[f] y) f y
----- λ
(*[f]) λy. f y
----- S ----- 0
(*[f] x) λy. f y    (*[f] x) x
----- @
(*[f] x) (λy. f y) x
----- λ
(*[f]) λx. (λy. f y) x    (*[f]) f
----- let
(*) let f = λx. (λy. f y) x
    in f
```

DFA: writing to file

minimised DFA: writing to file

```
spanning tree:
|F(L(A(S(F), 0)))
```

```
readback:
let F = λx. F x
in F
```



The term enriched by adbmals [18]. The adbmal (λ) is to be read as a scope delimiter that explicitly includes the name of the λ -variable whose scope it delimits. The adbmals are placed in accordance to the translation used.

A nameless scoped representation, where the names of abstraction variables are omitted for lambdas as well as for abstraction variable occurrences, shown as a 0-symbol. The scoping is expressed by scope-delimiters in the shape of an S-symbol. Note that these terms can be obtained by a simple syntactical transformation from the adbmals.

```
λ-letrec-term:
let f = λx. (λy. f x) x
in f
```

translation used: minimal prefix lengths

```
scoped (with adbmals):
let f = λx. (λy. /y. /x. f x) x
in f
```

```
scoped (with scope delimiters and nameless abstractions):
let f = λ. (λ. S((S(f) 0))) 0
in f
```



```
derivation:
(*[f]) f
----- S ----- 0
(*[f] x) f    (*[f] x) x
----- @
(*[f] x) f x
----- S
(*[f] x y) f x
----- λ ----- 0
(*[f] x) λy. f x    (*[f] x) x
----- @
(*[f] x) (λy. f x) x
----- λ
(*[f]) λx. (λy. f x) x    (*[f]) f
----- let
(*) let f = λx. (λy. f x) x
    in f
```

DFA: writing to file

minimised DFA: writing to file

```
spanning tree:
|F(L(A(L(S(A(S(F), |G(0))), G)))
```

```
readback:
let F = λx. let G = x
            in (λy. F G) G
in F
```



B.4 Figure 5

Also for the term $\lambda a. \lambda b. \text{let } f = a \text{ in } a a (f a) b$ from Fig. 5 on page 7 the translations are identical for $\llbracket \cdot \rrbracket_{\mathcal{H}}$ and $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$.

λ -letrec-term:

```
 $\lambda a. \lambda b. \text{let } f = a$ 
   $\text{in } a a (f a) b$ 
```

translation used: minimal prefix lengths

scoped (with adbmals):

```
 $\lambda a. \lambda b. \text{let } f = a$ 
   $\text{in } /b. a a (f a) b$ 
```

scoped (with scope delimiters and nameless abstractions):

```
 $\lambda. \lambda. \text{let } f = 0$ 
   $\text{in } S((0 0 (f 0))) 0$ 
```

derivation:

```

----- 0 ----- 0 ----- 0
(* a[f]) a  (* a[f]) a  (* a[f]) f  (* a[f]) a
----- @ ----- @
(* a[f]) a a  (* a[f]) f a
----- @
(* a[f]) a a (f a)
----- S ----- 0
(* a[f] b) a a (f a)  (* a[f] b) b
----- @
(* a[f]) a  (* a[f] b) a a (f a) b
----- let
(* a b) let f = a
  in a a (f a) b
----- λ
(* a) λb. let f = a
  in a a (f a) b
----- λ
(*) λa. λb. let f = a
  in a a (f a) b

```

DFA: writing to file



minimised DFA: writing to file

spanning tree:

```
L(L(A(S(A(|F(A(|G(0), G)), F)), 0)))
```

readback:

```
 $\lambda x. \text{let } F = G G$ 
   $G = x$ 
   $\text{in } \lambda y. F F y$ 
```



B.5 Example 5.14

For the terms L_1 , L_2 , and L_3 from Ex. 5.14 on page 11 the translations differ for $\llbracket \cdot \rrbracket_{\mathcal{H}}$ and $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$. Thus on the following pages we provide the output for both translations for each of the terms.

λ -letrec-term:

```

λx. let I = λz. z
    in λy. let f = x
          in y I (I y) (f f)

```

translation used: minimal prefix lengths

scoped (with adbmals):

```

λx. let I = λz. z
    in λy. let f = x
          in y /y. /x. I (/y. /x. I y) /y. f f

```

scoped (with scope delimiters and nameless abstractions):

```

λ. let I = λ. 0
  in λ. let f = 0
        in 0 S(S(I)) (S(S(I)) 0) S((f f))

```

derivation:

[illegible]

DFA: writing to file

```

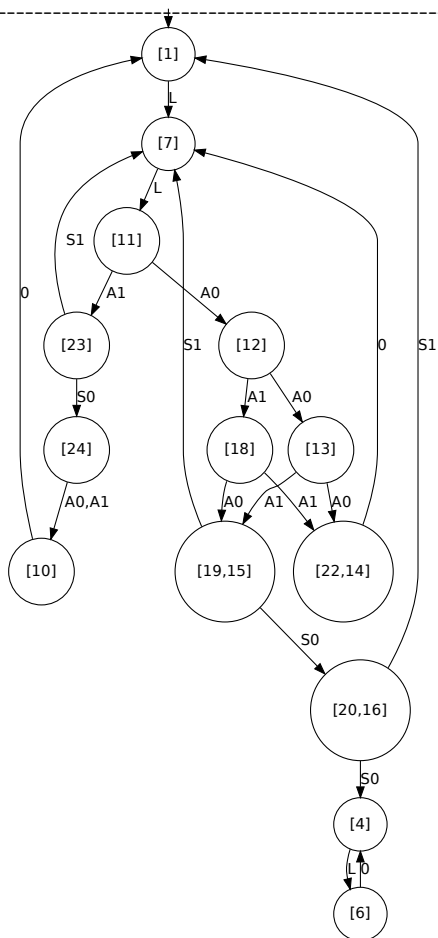
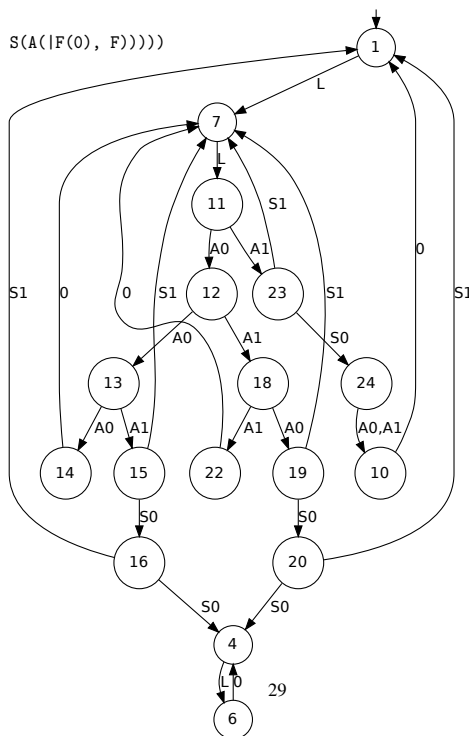
minimised DFA: writing to file

```

spanning tree:

$$L(L(A(A(A(|H(0), |G(S(S(L(0))))), A(G, H)), S(A(|F(0), F))))))$$

readback:

$$\lambda x. \text{ let } F = x$$
$$\quad \text{in } \lambda y. \text{ let } H = y$$
$$\quad \quad G = \lambda z. z$$
$$\quad \quad \text{in } H \ G \ (G \ H) \ (F \ F)$$



```

λ-letrec-term:
λx. let I = λz. z
  in λy. let f = x
    in λg = I
      in y g (g y) (f f)

```

translation used: maximal prefix lengths while maintaining eager scope-closure

```

scoped (with adbmals):
λx. let I = /x. λz. z
  in λy. let f = x
    in λg = /y. I
      in y g (g y) /y. f f

```

```

scoped (with scope delimiters and nameless abstractions):
λ. let I = S((λ. 0))
  in λ. let f = 0
    in λg = S(I)
      in 0 g (g 0) S((f f))

```

derivation:

$$\begin{array}{c}
 \begin{array}{c}
 \text{----- } 0 \\
 (* x[I f] y[g]) y \quad (* x[I f] y[g]) g \quad (* x[I f] y[g]) g \quad (* x[I f] y[g]) y \quad 0 \\
 \text{----- } @ \quad \text{----- } @ \\
 (* x[I f] y[g]) y g \quad (* x[I f] y[g]) g y \quad (* x[I f] y[g]) f f \\
 \text{----- } @ \quad \text{----- } @ \\
 (* x[I f] y[g]) y g (g y) \quad (* x[I f] y[g]) f f \\
 \text{----- } @ \\
 (* x[I f] y[g]) y g (g y) (f f)
 \end{array} \\
 \begin{array}{c}
 \text{----- } 0 \\
 (* x[I f]) I \\
 \text{----- } S \\
 (* x[I f]) x \quad (* x[I f] y[g]) I \quad (* x[I f] y[g]) y g (g y) (f f) \\
 \text{----- } 0 \\
 (* x[I] y) let f = x \\
 \quad g = I \\
 \quad in y g (g y) (f f)
 \end{array} \\
 \begin{array}{c}
 \text{----- } 0 \\
 (* z) z \\
 \text{----- } \lambda \\
 (*) \lambda z. z \quad (* x[I]) \lambda y. let f = x \\
 \quad g = I \\
 \quad in y g (g y) (f f)
 \end{array} \\
 \begin{array}{c}
 \text{----- } S \\
 (* x[I]) \lambda z. z \quad (* x[I]) \lambda y. let f = x \\
 \quad g = I \\
 \quad in y g (g y) (f f)
 \end{array} \\
 \begin{array}{c}
 (*) x) let I = \lambda z. z \\
 \quad in \lambda y. let f = x \\
 \quad \quad g = I \\
 \quad \quad in y g (g y) (f f)
 \end{array} \\
 \begin{array}{c}
 (*) \lambda x. let I = \lambda z. z \\
 \quad in \lambda y. let f = x \\
 \quad \quad g = I \\
 \quad \quad in y g (g y) (f f)
 \end{array}
 \end{array}$$

DFA: writing to file

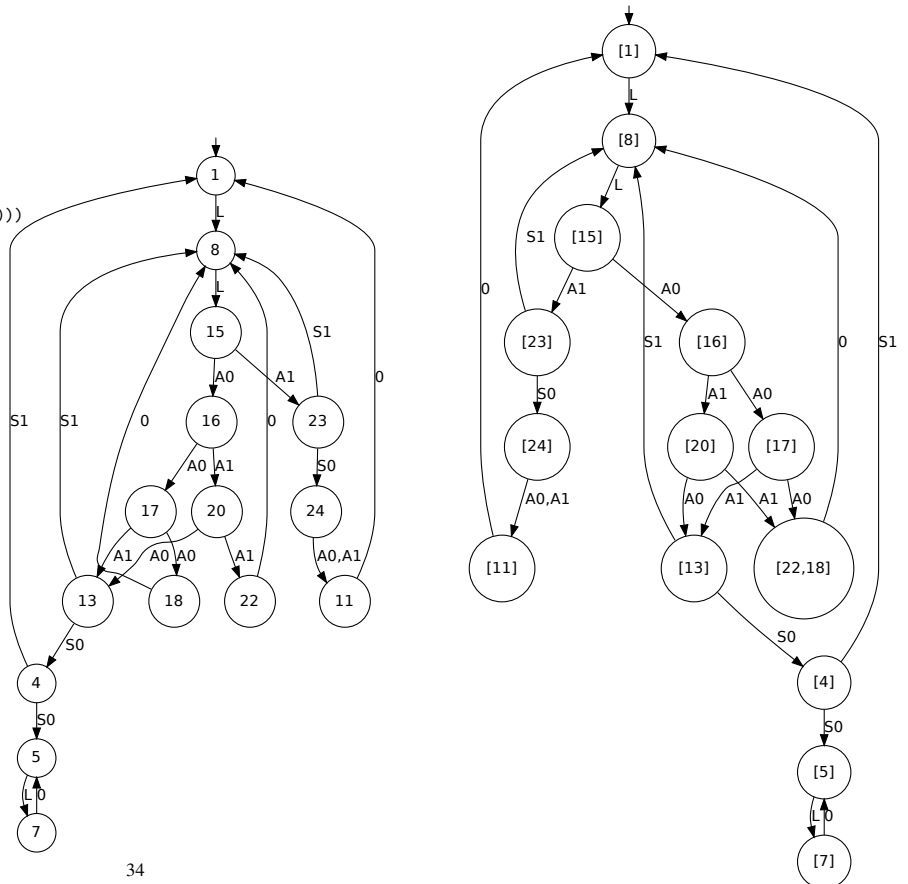
minimised DFA: writing to file

spanning tree:
L(L(A(A(IH(0), |G(S(S(L(0))))), A(G, H)), S(A(|F(0), F))))

```

readback:
λx. let F = x
  in λy. let H = y
    in λG = λz. z
      in H G (G H) (F F)

```



```
λ-letrec-term:
let I = λz. z
in λx. λy. let f = x
           in y I (I y) (f f)
```

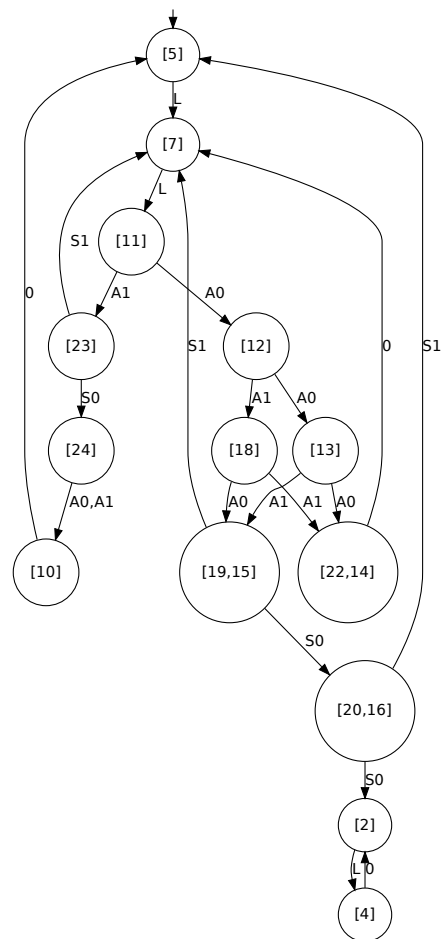
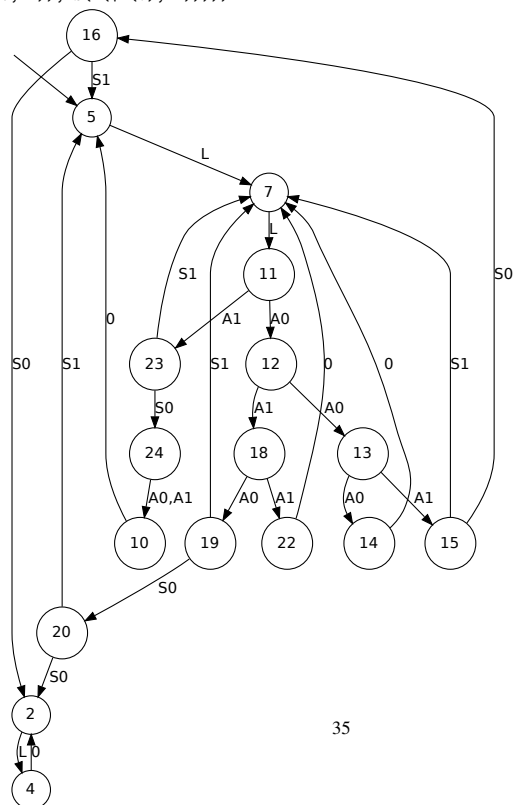
```
scoped (with adbmals):
let I = λz. z
in λx. λy. let f = x
           in y /y. /x. I (/y. /x. I y) /y. f f
```

derivation:

DFA: writing to file

spanning tree:
 $L(L(A(A(A(|H(0), |G(S(S(L(0))))), A(G, H)), S(A(|F(0), F))))))$

```
readback:
λx. let F = x
    in λy. let H = y
           G = λz. z
           in H G (G H) (F F)
```



B.6 Term L_2 from Example 5.14 (page 11)

This term has different translations for $\llbracket \cdot \rrbracket_{\mathcal{T}}$ and $\llbracket \cdot \rrbracket_{\mathcal{T}}^{\min}$. Thus on the following two pages we provide the output for both translations. See also, and compare with, the stepwise translations in Appendix A.

λ -letrec-term:

```

λx. λy. let I = λz. z
      f = x
      in y I (I y) (f f)

```

translation used: maximal prefix lengths while maintaining eager scope-closure

scoped (with adbmals):

```

λx. λy. let I = /y. /x. λz. z
      f = x
      in y I (I y) /y. f f

```

scoped (with scope delimiters and nameless abstractions):

```

λ. λ. let I = S(S((λ. 0)))
      f = 0
      in 0 I (I 0) S((f f))

```

derivation:

<pre> ----- 0 (* z) z ----- λ (*) λz. z ----- S (* x[f]) λz. z ----- S (* x[f] y[I]) λz. z ----- S (* x y) let I = λz. z f = x in y I (I y) (f f) ----- λ (* x) λy. let I = λz. z f = x in y I (I y) (f f) ----- λ (*) λx. λy. let I = λz. z f = x in y I (I y) (f f) </pre>	<pre> ----- 0 (* x[f] y[I]) y ----- 0 (* x[f] y[I]) I ----- 0 (* x[f] y[I]) I ----- 0 (* x[f] y[I]) y ----- 0 (* x[f] y[I]) y I ----- 0 (* x[f] y[I]) y I (I y) ----- 0 (* x[f] y[I]) y I (I y) (f f) ----- let </pre>	<pre> ----- 0 (* x[f]) f (* x[f]) f ----- 0 (* x[f]) f f ----- S (* x[f] y[I]) f f ----- 0 </pre>
--	--	---

DFA: writing to file

minimised DFA: writing to file

spanning tree:


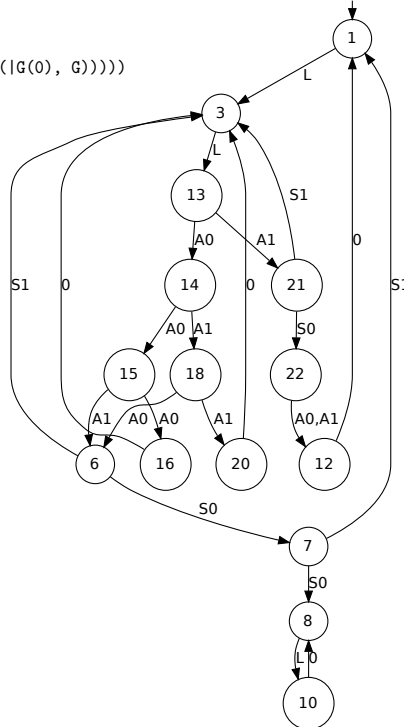
$L(L(A(A(|H(0), |F(S(S(L(0))))), A(F, H)), S(A(|G(0), G))))$

readback:

```

λx. let G = x
  in λy. let H = y
        F = λz. z
        in H F (F H) (G G)

```



```

λx. λy. let I = λz. z
      f = x
      in y I (I y) (f f)

```

scoped (with adbmals):

```

λx. λy. let I = λz. z
      f = x
      in y /y. /x. I (/y. /x. I y) /y. f f

```

scoped (with scope delimiters and nameless abstractions):

```

λ. λ. let I = λ. 0
      f = 0
      in 0 S(S(I)) (S(S(I)) 0) S((f f))

```

derivation:

$$\begin{array}{c}
 \text{(* [I] I)} \\
 \text{----- S} \\
 \text{(* [I] x [f]) I} \\
 \text{----- S} \\
 \text{(* [I] x [f] y) y} \quad \text{(* [I] x [f] y) I} \quad \text{(* [I] x [f] y) I} \quad \text{(* [I] x [f] y) y} \quad \text{(* [I] x [f]) f} \quad \text{(* [I] x [f]) f} \\
 \text{----- 0} \quad \text{----- S} \quad \text{----- S} \quad \text{----- 0} \quad \text{----- @} \quad \text{----- @} \\
 \text{(* [I] x [f] y) y I} \quad \text{(* [I] x [f] y) I y} \quad \text{(* [I] x [f] y) f f} \\
 \text{----- @} \quad \text{----- S} \\
 \text{(* [I] z) z} \quad \text{(* [I] x [f] y) y I (I y)} \quad \text{(* [I] x [f] y) f f} \\
 \text{----- 0} \quad \text{----- @} \\
 \text{(* [I]) } \lambda z. z \quad \text{(* [I] x [f]) x} \quad \text{(* [I] x [f] y) y I (I y) (f f)} \\
 \text{----- let} \\
 \text{(* x y) let I = } \lambda z. z \\
 \quad \text{f = x} \\
 \quad \text{in y I (I y) (f f)} \\
 \text{----- } \lambda \\
 \text{(* x) } \lambda y. \text{ let I = } \lambda z. z \\
 \quad \text{f = x} \\
 \quad \text{in y I (I y) (f f)} \\
 \text{----- } \lambda \\
 \text{(* } \lambda x. \lambda y. \text{ let I = } \lambda z. z \\
 \quad \text{f = x} \\
 \quad \text{in y I (I y) (f f)}
 \end{array}$$

DFA: writing to file

```

minimised DFA: writing to file

```

spanning tree:

$$L(L(A(A(|H(O), |G(S(S(L(O))))), A(G, H)), S(A(|F(O), F))))$$

readback:

```

λx. let F = x
    in λy. let H = y
           G = λz. z
           in H G (G H) (F F)

```

